BayRS Version 14.00

Part No. 308627-14.00 Rev 00
September 1999

4401 Great America Parkway
Santa Clara, CA 95054

Configuring IP, ARP, RARP, RIP, and OSPF Services

NORTEL NETWORKS™

Copyright © 1999 Nortel Networks

All rights reserved. Printed in the USA. September 1999.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel Networks NA Inc.

The software described in this document is furnished under a license agreement and may only be used in accordance with the terms of that license. A summary of the Software License is included in this document.

Trademarks

NORTEL NETWORKS is a trademark of Nortel Networks.

AN, BCN, FRE, and Bay Networks are registered trademarks and ASN, BayRS, BayStack, BCC, and System 5000 are trademarks of Nortel Networks.

All other trademarks and registered trademarks are the property of their respective owners.

Restricted Rights Legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, Nortel Networks NA Inc. reserves the right to make changes to the products described in this document without notice.

Nortel Networks NA Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product may be Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

Nortel Networks NA Inc. Software License Agreement

NOTICE: Please carefully read this license agreement before copying or using the accompanying software or installing the hardware unit with pre-enabled software (each of which is referred to as "Software" in this Agreement). BY COPYING OR USING THE SOFTWARE, YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. THE TERMS EXPRESSED IN THIS AGREEMENT ARE THE ONLY TERMS UNDER WHICH NORTEL NETWORKS WILL PERMIT YOU TO USE THE SOFTWARE. If you do not accept these terms and conditions, return the product, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

1. License Grant. Nortel Networks NA Inc. ("Nortel Networks") grants the end user of the Software ("Licensee") a personal, nonexclusive, nontransferable license: a) to use the Software either on a single computer or, if applicable, on a single authorized device identified by host ID, for which it was originally acquired; b) to copy the Software solely for backup purposes in support of authorized use of the Software; and c) to use and copy the associated user manual solely in support of authorized use of the Software by Licensee. This license applies to the Software only and does not extend to Nortel Networks Agent software or other Nortel Networks software products. Nortel Networks Agent software or other Nortel Networks software products are licensed for use under the terms of the applicable Nortel Networks NA Inc. Software License Agreement that accompanies such software and upon payment by the end user of the applicable license fees for such software.

2. Restrictions on use; reservation of rights. The Software and user manuals are protected under copyright laws. Nortel Networks and/or its licensors retain all title and ownership in both the Software and user manuals, including any revisions made by Nortel Networks or its licensors. The copyright notice must be reproduced and included with any copy of any portion of the Software or user manuals. Licensee may not modify, translate, decompile, disassemble, use for any competitive analysis, reverse engineer, distribute, or create derivative works from the Software or user manuals or any copy, in whole or in part. Except as expressly provided in this Agreement, Licensee may not copy or transfer the Software or user manuals, in whole or in part. The Software and user manuals embody Nortel Networks' and its licensors' confidential and proprietary intellectual property. Licensee shall not sublicense, assign, or otherwise disclose to any third party the Software, or any information about the operation, design, performance, or implementation of the Software and user manuals that is confidential to Nortel Networks and its licensors; however, Licensee may grant permission to its consultants, subcontractors, and agents to use the Software at Licensee's facility, provided they have agreed to use the Software only in accordance with the terms of this license.

3. Limited warranty. Nortel Networks warrants each item of Software, as delivered by Nortel Networks and properly installed and operated on Nortel Networks hardware or other equipment it is originally licensed for, to function substantially as described in its accompanying user manual during its warranty period, which begins on the date Software is first shipped to Licensee. If any item of Software fails to so function during its warranty period, as the sole remedy Nortel Networks will at its discretion provide a suitable fix, patch, or workaround for the problem that may be included in a future Software release. Nortel Networks further warrants to Licensee that the media on which the Software is provided will be free from defects in materials and workmanship under normal use for a period of 90 days from the date Software is first shipped to Licensee. Nortel Networks will replace defective media at no charge if it is returned to Nortel Networks during the warranty period along with proof of the date of shipment. This warranty does not apply if the media has been damaged as a result of accident, misuse, or abuse. The Licensee assumes all responsibility for selection of the Software to achieve Licensee's intended results and for the installation, use, and results obtained from the Software. Nortel Networks does not warrant a) that the functions contained in the software will meet the Licensee's requirements, b) that the Software will operate in the hardware or software combinations that the Licensee may select, c) that the operation of the Software will be uninterrupted or error free, or d) that all defects in the operation of the Software will be corrected. Nortel Networks is not obligated to remedy any Software defect that cannot be reproduced with the latest Software release. These warranties do not apply to the Software if it has been (i) altered, except by Nortel Networks or in accordance with its instructions; (ii) used in conjunction with another vendor's product, resulting in the defect; or (iii) damaged by improper environment, abuse, misuse, accident, or negligence. THE FOREGOING WARRANTIES AND LIMITATIONS ARE EXCLUSIVE REMEDIES AND ARE IN LIEU OF ALL OTHER WARRANTIES EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Licensee is responsible

for the security of its own data and information and for maintaining adequate procedures apart from the Software to reconstruct lost or altered files, data, or programs.

4. Limitation of liability. IN NO EVENT WILL NORTEL NETWORKS OR ITS LICENSORS BE LIABLE FOR ANY COST OF SUBSTITUTE PROCUREMENT; SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES; OR ANY DAMAGES RESULTING FROM INACCURATE OR LOST DATA OR LOSS OF USE OR PROFITS ARISING OUT OF OR IN CONNECTION WITH THE PERFORMANCE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF NORTEL NETWORKS RELATING TO THE SOFTWARE OR THIS AGREEMENT EXCEED THE PRICE PAID TO NORTEL NETWORKS FOR THE SOFTWARE LICENSE.

5. Government Licensees. This provision applies to all Software and documentation acquired directly or indirectly by or on behalf of the United States Government. The Software and documentation are commercial products, licensed on the open market at market prices, and were developed entirely at private expense and without the use of any U.S. Government funds. The license to the U.S. Government is granted only with restricted rights, and use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1) of the Commercial Computer Software—Restricted Rights clause of FAR 52.227-19 and the limitations set out in this license for civilian agencies, and subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause of DFARS 252.227-7013, for agencies of the Department of Defense or their successors, whichever is applicable.

6. Use of Software in the European Community. This provision applies to all Software acquired for use within the European Community. If Licensee uses the Software within a country in the European Community, the Software Directive enacted by the Council of European Communities Directive dated 14 May, 1991, will apply to the examination of the Software to facilitate interoperability. Licensee agrees to notify Nortel Networks of any such intended examination of the Software and may procure support and assistance from Nortel Networks.

7. Term and termination. This license is effective until terminated; however, all of the restrictions with respect to Nortel Networks' copyright in the Software and user manuals will cease being effective at the date of expiration of the Nortel Networks copyright; those restrictions relating to use and disclosure of Nortel Networks' confidential information shall continue in effect. Licensee may terminate this license at any time. The license will automatically terminate if Licensee fails to comply with any of the terms and conditions of the license. Upon termination for any reason, Licensee will immediately destroy or return to Nortel Networks the Software, user manuals, and all copies. Nortel Networks is not liable to Licensee for damages in any form solely by reason of the termination of this license.

8. Export and Re-export. Licensee agrees not to export, directly or indirectly, the Software or related technical data or information without first obtaining any required export licenses or other governmental approvals. Without limiting the foregoing, Licensee, on behalf of itself and its subsidiaries and affiliates, agrees that it will not, without first obtaining all export licenses and approvals required by the U.S. Government: (i) export, re-export, transfer, or divert any such Software or technical data, or any direct product thereof, to any country to which such exports or re-exports are restricted or embargoed under United States export control laws and regulations, or to any national or resident of such restricted or embargoed countries; or (ii) provide the Software or related technical data or information to any military end user or for any military end use, including the design, development, or production of any chemical, nuclear, or biological weapons.

9. General. If any provision of this Agreement is held to be invalid or unenforceable by a court of competent jurisdiction, the remainder of the provisions of this Agreement shall remain in full force and effect. This Agreement will be governed by the laws of the state of California.

Should you have any questions concerning this Agreement, contact Nortel Networks, 4401 Great America Parkway, P.O. Box 58185, Santa Clara, California 95054-8185.

LICENSEE ACKNOWLEDGES THAT LICENSEE HAS READ THIS AGREEMENT, UNDERSTANDS IT, AND AGREES TO BE BOUND BY ITS TERMS AND CONDITIONS. LICENSEE FURTHER AGREES THAT THIS AGREEMENT IS THE ENTIRE AND EXCLUSIVE AGREEMENT BETWEEN NORTEL NETWORKS AND LICENSEE, WHICH SUPERSEDES ALL PRIOR ORAL AND WRITTEN AGREEMENTS AND COMMUNICATIONS BETWEEN THE PARTIES PERTAINING TO THE SUBJECT MATTER OF THIS AGREEMENT. NO DIFFERENT OR ADDITIONAL TERMS WILL BE ENFORCEABLE AGAINST NORTEL NETWORKS UNLESS NORTEL NETWORKS GIVES ITS EXPRESS WRITTEN CONSENT, INCLUDING AN EXPRESS WAIVER OF THE TERMS OF THIS AGREEMENT.

Contents

Preface

Before You Begin xvii
Text Conventions xviii
Acronyms xix
Related Publications xxi
How to Get Help xxii

Chapter 1
IP Concepts, Terminology, and Features

IP Addresses 1-2
Subnet Addressing 1-4
Supernet Addressing 1-7
Classless Interdomain Routing 1-8
Autonomous Systems 1-8
Address Resolution Protocol (ARP) 1-10
Reverse Address Resolution Protocol (RARP) 1-12
Routing Information Protocol (RIP) 1-13
Open Shortest Path First (OSPF) Protocol 1-13
OSPF Addresses and Variable-Length Masks 1-14
OSPF Neighbors 1-14
Neighbor Adjacencies 1-14
Designated Routers 1-15
OSPF Areas 1-16
Stub Area 1-17
Not-So-Stubby Area (NSSA) 1-17
OSPF Router Types 1-18
AS External Routes 1-18
OSPF Implementation Notes 1-19
Router Discovery Protocol 1-20

Route Preferences 1-20
Route Weights 1-21
IP Routing Policies and Filters 1-22
How OSPF Implements Routing Policies 1-25
IP Traffic Filters 1-26

Chapter 2
Starting IP Services

Using the BCC 2-2
Starting IP 2-2
Step 1: Configuring a Physical Interface 2-2
Step 2: Configuring an IP Interface 2-3
Starting RIP 2-3
Starting OSPF 2-4
Starting Router Discovery 2-4
Using Site Manager 2-5
Starting IP 2-5
Starting RIP 2-6
Adding RIP to an IP Interface 2-7
Starting OSPF 2-8
Starting RARP 2-9
Configuring a Circuitless IP Interface 2-10
Choosing Slots to Support the Circuitless Interface 2-11
Configuring an Unnumbered IP Interface 2-11
Using the Alternate Associated Address Option 2-13

Chapter 3
Configuring and Customizing IP

Customizing IP Global Parameters 3-2
Disabling and Reenabling Global IP 3-3
Configuring the Router for Not-Forwarding Mode 3-4
Configuring Bridging on a Router in Not-Forwarding Mode 3-6
Setting the Time-to-Live Value on a Source Packet 3-9
Allowing an All-Zero or All-One Subnet Address 3-10
Estimating the Size of the Routing Table 3-11
Using a Default Route for an Unknown Subnet 3-12

Specifying the Maximum Number of IP Policies 3-13
Disabling and Reenabling Route Filter Support 3-14
Configuring Equal-Cost Multipath Support for RIP and OSPF 3-15
Customizing the IP Routing Table Structure 3-18
Specifying the Percentage of Buffers Available to ARP 3-19
Customizing an IP Interface 3-20
Disabling and Reenabling an IP Interface 3-22
Specifying a Broadcast Address for an Interface 3-23
Specifying the Cost of an Interface 3-24
Enabling MTU Discovery on an Interface 3-26
Enabling and Disabling ICMP Address-Mask Replies 3-27
Disabling and Reenabling ICMP Redirect Messages 3-29
Enabling All-Subnet Broadcasting on an Interface 3-30
Disabling UDP Checksum Processing on an Interface 3-31
Specifying a MAC Address or E.164 Address 3-33
Enabling Source Routing over a Token Ring Network 3-34
Configuring the MTU for an IP Interface 3-37
Configuring an SMDS Address 3-38
Configuring a WAN Address for a Frame Relay Network 3-39
Specifying the Maximum Size of the Forwarding Table 3-40
Configuring an Interface for an ATM Logical IP Subnet 3-42
Deleting IP from an Interface 3-45
Configuring an Adjacent Host Address 3-46
Defining a Static Route 3-49
Defining a Static Default Route 3-53
Defining a Static Black Hole for a Supernet 3-53
Configuring and Customizing Router Discovery 3-54
Enabling and Disabling Router Discovery 3-55
Choosing a Broadcast Type 3-56
Specifying the Time Interval Between Advertisements 3-57
Configuring the Lifetime of Advertised Addresses 3-58
Specifying Interface Preference 3-59

Chapter 4
Customizing ARP and RARP Services

Customizing ARP 4-2
Enabling and Disabling Global ARP 4-2
Customizing Global ARP 4-3
Selecting an Address Resolution Scheme for an IP Interface 4-5
Selecting an Encapsulation Option for ARP and Probe 4-7
Enabling Proxy ARP on an Interface 4-8
Timing Out Entries in the Address Resolution Cache 4-10
Customizing RARP 4-12
Disabling and Enabling RARP on an Interface 4-13
Defining the RARP Mapping Table 4-13
Disabling and Reenabling RARP Globally 4-14
Deleting RARP Globally 4-15

Chapter 5
Customizing RIP Services

Setting the RIP Diameter 5-2
Customizing a RIP Interface 5-3
Disabling and Reenabling RIP on an Interface 5-3
Selecting the RIP Version 5-4
Authenticating the Password on a Version 2 Update 5-6
Supplying RIP Updates on an Interface 5-8
Specifying the Update Mode 5-9
Sending Triggered Updates 5-11
Specifying a Time-to-Live Value 5-12
Receiving RIP Updates on an Interface 5-14
Supplying a Default Route on an Interface 5-15
Listening for a Default Route 5-17
Configuring a RIP Interface for Dial-Optimized Routing 5-18
Specifying an Update Interval 5-18
Specifying a Timeout Period 5-20
Specifying a Holddown Period 5-21
Specifying a Stabilization Time 5-22
Deleting RIP from an Interface 5-23

Configuring RIP Accept and Announce Policies 5-24
Defining a RIP Accept Policy 5-25
Supplying Modification Values for a RIP Accept Policy 5-28
Specifying Matching Criteria for a RIP Accept Policy 5-29
Defining a RIP Announce Policy 5-31
Supplying Modification Values for a RIP Announce Policy 5-33
Specifying Matching Criteria for a RIP Announce Policy 5-35

Chapter 6
Customizing OSPF Services

Customizing OSPF Global Parameters 6-2
Enabling and Disabling OSPF on the Router 6-3
Supplying an OSPF Router ID 6-4
Configuring the Soloist and Backup Soloist on a Slot 6-5
Enabling the Boundary Function 6-7
Configuring the Metric Type for an ASE Advertisement 6-8
Choosing a Tag Generation Method for an ASE Advertisement 6-11
Setting the Holddown Timer 6-14
Configuring Message Logging 6-15
Configuring External Route Preference 6-17
Customizing OSPF on an IP Interface 6-18
Enabling and Disabling OSPF on an Interface 6-19
Configuring an Area ID 6-20
Specifying the Interface Type 6-21
Specifying Router Priority for a Multiaccess Network 6-23
Estimating the Transit Delay 6-24
Setting the Retransmit Interval 6-25
Setting the Hello Interval 6-27
Setting the Dead Interval 6-29
Setting the Poll Interval for NBMA Neighbors 6-30
Specifying the Metric Cost 6-31
Specifying the MTU Size 6-34
Specifying the MTU Mismatch Detect 6-35
Configuring a Neighbor on an NBMA Interface 6-36
Deleting OSPF from an IP Interface 6-38

Defining an Area 6-40
Supplying an ID for the Area 6-40
Disabling and Reenabling an Area 6-42
Configuring Authentication 6-43
Configuring a Summary Route 6-44
Configuring a Stub Area 6-46
Configuring a Not-So-Stubby Area (NSSA) 6-48
Configuring an NSSA Address Range 6-50
Configuring an Area Border Router 6-53
Configuring a Virtual Backbone Link Through a Transit Area 6-54
Configuring OSPF Accept and Announce Policies 6-57
Defining an OSPF Accept Policy 6-58
Specifying Matching Criteria for an OSPF Accept Policy 6-61
Defining an OSPF Announce Policy 6-63
Specifying Modification Values for an OSPF Announce Policy 6-65
Specifying Matching Criteria for an OSPF Announce Policy 6-68

Appendix A
Site Manager Parameters

IP Parameters A-2
IP Configuration Parameters A-2
IP Interface Parameters A-4
IP Global Parameters A-19
Static Route Parameters A-26
Adjacent Host Parameters A-30
Router Discovery Parameters A-33
RARP Parameters A-35
OSPF Parameters A-37
OSPF Global Parameters A-37
OSPF Interface Parameters A-42
Neighbor Parameters for an NBMA Interface A-49
OSPF Area Parameters A-51
Area Range Parameters A-55
OSPF NSSA Range Parameters A-58
OSPF Virtual Interface Parameters A-61
RIP Parameters A-64

Appendix B
Routing Policy Parameters

Common Accept Policy Parameters B-2
RIP Accept Policy Parameters B-5
OSPF Accept Policy Parameters B-6
Common Announce Policy Parameters B-7
RIP Announce Policy Parameters B-17
OSPF Announce Policy Parameters B-19

Appendix C
Import and Export Route Filters

RIP Import Filters C-1
RIP Export Filters C-5
OSPF Import Filters C-8
OSPF Export Filters C-9

Appendix D
Route Weight Worksheet

Appendix E
IP/OSPF Configuration Example

Index

Figures

Figure 1-1. Network and Host Portions of IP Addresses 1-3
Figure 1-2. Supernet Addresses 1-7
Figure 1-3. Internetwork Segmented into Three Autonomous Systems 1-9
Figure 1-4. ARP Example 1-10
Figure 1-5. RARP Server Supplying an IP Address 1-12
Figure 1-6. OSPF Areas 1-16
Figure 1-7. IP Routing Table 1-23
Figure 1-8. Accept and Announce Policies 1-24
Figure 3-1. IP Interface 3-20
Figure 3-2. IP Routers Source Routing Across a Token Ring Network 3-35
Figure 4-1. Proxy ARP Example 4-8
Figure 6-1. OSPF ASE Routes 6-9
Figure 6-2. AS External Route Tag 6-12
Figure 6-3. Example of Using Configurable Cost Metrics 6-32
Figure 6-4. Area Border Router 6-53
Figure 6-5. Virtual Link and Transit Area 6-55
Figure E-1. IP/OSPF Configuration E-2

Tables

Table 1-1. Network Classes and IP Addresses 1-2
Table 1-2. Subnet Masks for Class B and Class C Addresses 1-6
Table 3-1. Source Routing Bridge Support for Host-Only Mode 3-6
Table 3-2. Learning Bridge Support for Host-Only Mode 3-7
Table 3-3. MAC Address Parameter Settings 3-33
Table 3-4. BCC Adjacent Host Parameters 3-47
Table 3-5. BCC Static Route Parameters 3-51
Table 4-1. BCC Global ARP Parameters 4-3
Table 5-1. BCC Definition Parameters for RIP Accept Policies 5-26
Table 5-2. BCC Modification Parameter for RIP Accept Policies 5-28
Table 5-3. BCC Match Parameters for RIP Accept Policies 5-30
Table 5-4. BCC Definition Parameters for RIP Announce Policies 5-32
Table 5-5. BCC Modification Parameters for RIP Announce Policies 5-34
Table 5-6. BCC Match Parameters for RIP Announce Policies 5-35
Table 6-1. OSPF Log Messages 6-15
Table 6-2. OSPF Interface Types 6-22
Table 6-3. Retransmit Interval Settings 6-25
Table 6-4. Hello Interval Settings 6-27
Table 6-5. Dead Interval Settings 6-29
Table 6-6. Cost Settings 6-32
Table 6-7. BCC Neighbor Parameters for an NBMA Interface 6-37
Table 6-8. BCC Parameters for OSPF NSSA 6-49
Table 6-9. BCC Address Range Parameters for OSPF NSSA 6-51
Table 6-10. BCC Definition Parameters for OSPF Accept Policies 6-59
Table 6-11. BCC Match Parameters for OSPF Accept Policies 6-61
Table 6-12. BCC Definition Parameters for OSPF Announce Policies 6-64
Table 6-13. BCC Modification Parameters for OSPF Announce Policies 6-66
Table 6-14. BCC Match Parameters for OSPF Announce Policies 6-68
Table E-1. Internal Backbone Router 1 ..... E-3

Table E-2. Area Border Router 2 ..... E-4
Table E-3. Area Border Router 3 ..... E-5
Table E-4. Area Border Router 4 ..... E-6
Table E-5. Internal Backbone Router 5 ..... E-7
Table E-6. AS Boundary Router 6 ..... E-7

Preface

This guide describes Internet Protocol (IP) services and what you do to start and customize IP services on a Nortel Networks™ router.

You can use the Bay Command Console (BCC™) or Site Manager to configure IP services on a router. In this guide, you will find instructions for using both the BCC and Site Manager.

Before You Begin

Before using this guide, you must complete the following procedures. For a new router:

- Install the router (see the installation guide that came with your router).
- Connect the router to the network and create a pilot configuration file (see *Quick-Starting Routers*, *Configuring BayStack Remote Access*, or *Connecting ASN Routers to a Network)*.

Make sure that you are running the latest version of Nortel Networks BayRS™ and Site Manager software. For information about upgrading BayRS and Site Manager, see the upgrading guide for your version of BayRS.

Text Conventions

This guide uses the following text conventions:

angle brackets (< >)	Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command. Example: If the command syntax is: **ping** <*ip_address*>, you enter: **ping 192.32.10.12**
bold text	Indicates command names and options and text that you need to enter. Example: Enter **show ip** {**alerts** | **routes**}. Example: Use the **dinfo** command.
braces ({})	Indicate required elements in syntax descriptions where there is more than one option. You must choose only one of the options. Do not type the braces when entering the command. Example: If the command syntax is: **show ip** {**alerts** | **routes**}, you must enter either: **show ip alerts** or **show ip routes**, but not both.
brackets ([ ])	Indicate optional elements in syntax descriptions. Do not type the brackets when entering the command. Example: If the command syntax is: **show ip interfaces** [**-alerts**], you can enter either: **show ip interfaces** or **show ip interfaces -alerts**.
ellipsis points (. . .)	Indicate that you repeat the last element of the command as needed. Example: If the command syntax is: **ethernet/2/1** [<*parameter*> <*value*>] **. . .** , you enter **ethernet/2/1** and as many parameter-value pairs as needed.

italic text	Indicates file and directory names, new terms, book titles, and variables in command syntax descriptions. Where a variable is two or more words, the words are connected by an underscore. Example: If the command syntax is: **show at** <*valid_route*> *valid_route* is one variable and you substitute one value for it.
`screen text`	Indicates system output, for example, prompts and system messages. Example: `Set Trap Monitor Filters`
separator (>)	Shows menu paths. Example: Protocols > IP identifies the IP option on the Protocols menu.
vertical line (|)	Separates choices for command keywords and arguments. Enter only one of the choices. Do not type the vertical line when entering the command. Example: If the command syntax is: **show ip** {**alerts** | **routes**}, you enter either: **show ip alerts** or **show ip routes**, but not both.

Acronyms

This guide uses the following acronyms:

ARE	all routes explorer
ARP	Address Resolution Protocol
AS	autonomous system
ASB	all-subnet broadcast
ASBR	AS boundary router
ASE	autonomous system external
ATM	asynchronous transfer mode

BGP	Border Gateway Protocol
BR	border router
CIDR	classless interdomain routing
DVMRP	Distance Vector Multicast Routing Protocol
ECMP	equal-cost multipath
EGP	Exterior Gateway Protocol
FDDI	Fiber Distributed Data Interface
FIFO	first in first out
ICMP	Internet Control Message Protocol
IGP	interior gateway protocol
IP	Internet Protocol
ISP	Internet service provider
LIS	logical IP subnet
LSA	link state advertisement
LSDB	link state database
MAC	media access control
MOSPF	Multicasting Open Shortest Path First
MTU	maximum transmission unit
NBMA	nonbroadcast multi-access
NIC	Network Information Center
NSSA	not-so-stubby area
OSPF	Open Shortest Path First
PPP	Point-to-Point Protocol
RARP	Reverse Address Resolution Protocol
RFC	Request for Comments
RIF	routing information field
RIP	Routing Information Protocol
SMDS	Switched Multimegabit Data Service
STE	spanning tree explorer

TTL	time to live
UDP	User Datagram Protocol
VC	virtual circuit
WAN	wide area network

Related Publications

For more information about IP services, refer to the following publications:

- *Reference for BCC IP show Commands* (Bay Networks part number 308603-14.00 Rev 00)

  Provides descriptions of all **show** commands for IP services, including the commands that display IP, ARP, RIP, and OSPF configuration and statistical data.

- *Configuring IP Exterior Gateway Protocols (BGP and EGP)* (Bay Networks part number 308628-14.00 Rev 00)

  Provides a description of Border Gateway Protocol (BGP) and Exterior Gateway Protocol (EGP) services and instructions for configuring them.

- *Configuring GRE, NAT, RIPSO, and BFE Services* (Bay Networks part number 308625-14.00 Rev 00)

  Provides a description of Generic Routing Encapsulation (GRE), Network Address Translation (NAT), Revised IP Security Option (RIPSO), and Blacker front-end services and instructions for configuring them.

You can print selected technical manuals and release notes free, directly from the Internet. Go to *support.baynetworks.com/library/tpubs/*. Find the product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Using Adobe Acrobat Reader, you can open the manuals and release notes, search for the sections you need, and print them on most standard printers. You can download Acrobat Reader free from the Adobe Systems Web site, *www.adobe.com*.

You can purchase selected documentation sets, CDs, and technical publications through the collateral catalog. The catalog is located on the World Wide Web at *support.baynetworks.com/catalog.html* and is divided into sections arranged alphabetically:

- The “CD ROMs” section lists available CDs.
- The “Guides/Books” section lists books on technical topics.
- The “Technical Manuals” section lists available printed documentation sets.

How to Get Help

If you purchased a service contract for your Nortel Networks product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Nortel Networks service program, contact one of the following Nortel Networks Technical Solutions Centers:

Technical Solutions Center	Telephone Number
Billerica, MA	800-2LANWAN (800-252-6926)
Santa Clara, CA	800-2LANWAN (800-252-6926)
Valbonne, France	33-4-92-96-69-68
Sydney, Australia	61-2-9927-8800
Tokyo, Japan	81-3-5402-7041

Chapter 1
IP Concepts, Terminology, and Features

This chapter introduces concepts and terminology used in this guide.

Topic	Page
IP Addresses	1-2
Autonomous Systems	1-8
Address Resolution Protocol (ARP)	1-10
Reverse Address Resolution Protocol (RARP)	1-12
Routing Information Protocol (RIP)	1-13
Open Shortest Path First (OSPF) Protocol	1-13
Router Discovery Protocol	1-20
Route Preferences	1-20
Route Weights	1-21
IP Routing Policies and Filters	1-22
IP Traffic Filters	1-26

IP Addresses

An IP address consists of 32 bits that have the form *network.host*. The *network* portion is a network number ranging from 8 to 24 bits. The *host* portion is the remaining 8 to 24 bits identifying a specific host on the network. The Internet Network Information Center (NIC) assigns the network portion of the IP address. Your network administrator assigns the host portion.

NIC recognizes three primary classes of networks: A, B, and C. In addition, NIC has recently identified two other classes: Class D for networks that support multicasting, which allows an IP datagram to be transmitted to a single multicast group consisting of hosts spread across separate physical networks; and Class E for experimental networks. The IP router does not fully support Class D or Class E networks.

Based on the size of the network, NIC classifies a network as Class A, B, or C (the most common). The network class determines the number of bits assigned to the network and host portions of the IP address (Table 1-1).

Table 1-1. Network Classes and IP Addresses

Network Size	Class	Network Portion	Host Portion
More than 65,534 hosts	A	8 bits	24 bits
254 to 65,533 hosts	B	16 bits	16 bits
Fewer than 254 hosts	C	24 bits	8 bits

The position of the first bit set to 0 (whether it is the first, second, third, or fourth bit) in the first octet of an IP address indicates the network class (A, B, C, or D). If no bit is set to 0, it is a Class E network. Figure 1-1 shows the placement of the first bit set to 0 for Class A, B, and C networks. The figure also shows how a network's class affects the network and host portions of the IP address.

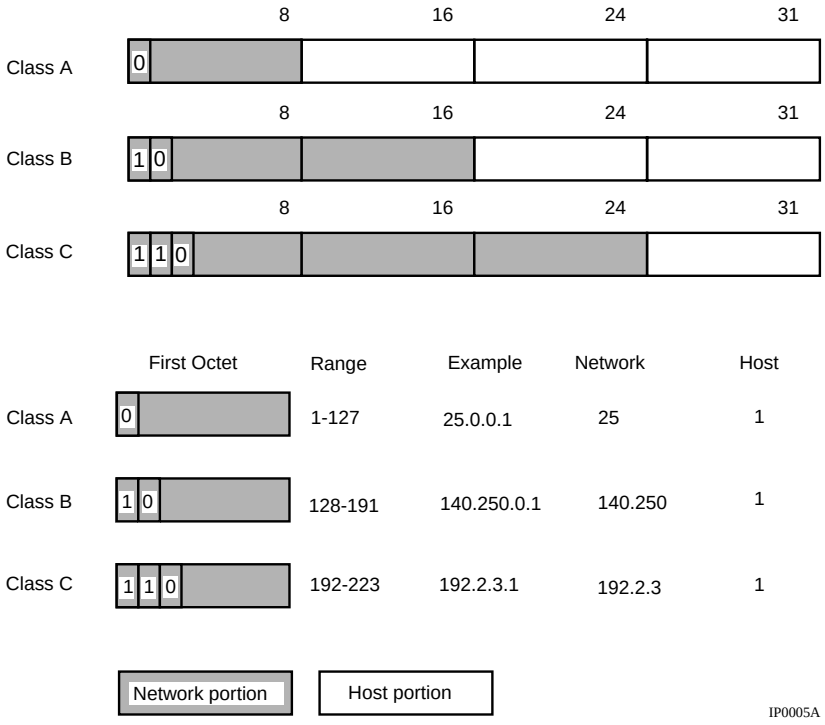

Figure 1-1. Network and Host Portions of IP Addresses

You specify IP addresses in dotted-decimal notation. To express an IP address in dotted-decimal notation, you convert each 8-bit octet of the IP address to a decimal number and separate the numbers by decimal points.

For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167. The most significant 2 bits (10) in the first octet indicate that the network is Class B; therefore, the first 16 bits compose the NIC-assigned network portion field. The third octet (00001010) and fourth octet (10100111) compose the host field.

Subnet Addressing

The concept of subnetworks (or subnets) extends the IP addressing scheme. Subnets are two or more physical networks that share a common network-identification field (the NIC-assigned network portion of the 32-bit IP address). Subnets allow an IP router to hide the complexity of multiple LANs from the rest of the internetwork.

Using subnets, you partition the host portion of an IP address into a subnet number and a "real" host number on that subnet. The IP address is then defined by *network.subnet.host*. Routers outside the network do not interpret the subnet and host portions of the IP address separately.

Routers within a network containing subnets use a 32-bit subnet mask that identifies the extension bits. In *network.subnet.host*, the *subnet.host* portion (or the local portion) contains an arbitrary number of bits. The network administrator allocates bits within the local portion to subnet and host, and then assigns values to subnet and host.

For example, the following is the IP address of a network that contains subnets: 10000000 00100000 00001010 10100111. You specify this address in dotted-decimal notation as 128.32.10.167.

The second bit of the first octet is set to 0, indicating that the network is a Class B network. Therefore, the NIC-assigned network portion contains 16 bits, and the locally assigned local portion contains 16 bits.

The network administrator allocates the 16 bits in the local portion field as follows:

- Upper 8 bits (00001010) with a value of 10 to the subnet portion
- Lower 8 bits (10100111) with a value of 167 to the host portion

In other words, the 16-bit local portion field, together with the 16-bit network field, specify host 167 on subnet 10 of network 128.32.

You now need a subnet mask to identify those bits in the 32-bit IP address that specify the network field and those bits that specify the subnet field. Like the IP address, you specify the subnet mask in dotted-decimal notation.

You construct a subnet mask as follows:

- Assign a value of 1 to each of the 8, 16, or 24 bits in the network field.
- Assign a value of 1 to each bit in the subnet field.
- Assign a value of 0 to each bit in the host field.
- Convert the resulting 32-bit string to dotted-decimal notation.

For example, to construct a subnet mask for the IP address 10000000 00100000 00001010 10100111, do the following:

1. Assign a value of 1 to each bit in the network field.

   The position of the first bit set to 0 in the first octet of the IP address indicates that the network is Class B; therefore, the network field contains 16 bits: 11111111 11111111.

2. Assign a value of 1 to each bit in the subnet field.

   Allocate the upper 8 bits of the local portion to the subnet portion, as follows: 11111111.

3. Assign a value of 0 to each bit in the host field.

   Allocate the lower 8 bits of the local portion field to the host identification, as follows: 00000000.

4. Convert the resulting 32-bit string (11111111 11111111 11111111 00000000) to dotted-decimal notation, as follows: 255.255.255.000.

Table 1-2 lists the range of possible subnet masks for Class B and Class C addresses, along with the number of bits that the mask allocates for a subnet address, the number of recommended subnets associated with the mask, and the number of hosts per subnet.

Table 1-2. Subnet Masks for Class B and Class C Addresses

Number of Bits	Subnet Mask	Number of Subnets (Recommended)	Number of Hosts per Subnet
Class B			
2	255.255.192.0	2	16,382
3	255.255.224.0	6	8,190
4	255.255.240.0	14	4,094
5	255.255.248.0	30	2,046
6	255.255.252.0	62	1,022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1,022	62
11	255.255.255.224	2,046	30
12	255.255.255.240	4,094	14
13	255.255.255.248	8,190	6
14	255.255.255.252	16,382	2
Class C			
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

Supernet Addressing

A supernet is a group of networks identified by contiguous network addresses. Internet service providers (ISPs) can assign customers blocks of contiguous addresses to define supernets as needed.

Each supernet has a unique supernet address that consists of the upper bits shared by all addresses in the contiguous block. For example, Figure 1-2 illustrates a block of contiguous 32-bit addresses (192.32.0.0 through 192.32.7.0 in dotted-decimal notation).

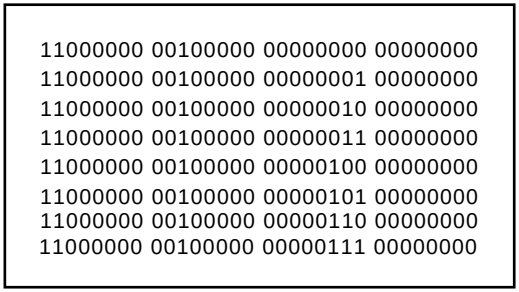

Figure 1-2. Supernet Addresses

The supernet address for this block is 11000000 00100000 00000, the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an *address*/*mask* pair:

- *address* is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in dotted-decimal notation).
- *mask* is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 11111000 00000000 (255.255.248.0 in dotted-decimal notation).

The complete supernet address in this example is 192.32.0.0/255.255.248.0.

Classless Interdomain Routing

Classless interdomain routing (CIDR) is an addressing scheme that employs supernet addresses to represent multiple IP destinations. Rather than advertise a separate route for each destination in a supernet, a router can use a supernet address to advertise a single route -- called an *aggregate* route -- that represents all of the destinations. CIDR reduces the size of the routing tables used to store advertised IP routes.

OSPF supports classless routing within a domain.

Autonomous Systems

LANs and WANs interconnected by IP routers form a group of networks called an *internetwork*. For administrative purposes, an internetwork is divided into autonomous systems. An *autonomous system* (AS) is simply a collection of routers (called *gateways* in IP terminology) and hosts. Figure 1-3 depicts a sample internetwork segmented into three autonomous systems.

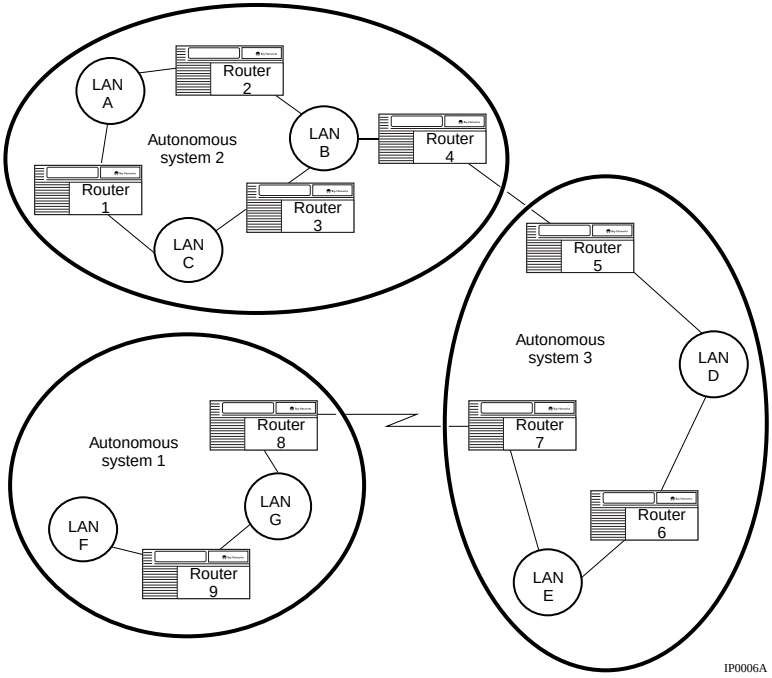

Figure 1-3. Internetwork Segmented into Three Autonomous Systems

Address Resolution Protocol (ARP)

The IP router needs both a physical address and an IP address to transmit a datagram. If the router knows only the network host's IP address, the Address Resolution Protocol (ARP) enables the router to determine the network host's physical address by binding a 32-bit IP address to a 48-bit media access control (MAC) address. A router can use ARP across a single network only, and the network hardware must support physical broadcasts.

For example, in Figure 1-4, the router and host C are on the same physical network. Both devices have an assigned IP address (the router's is 140.250.200.1 and host C's is 140.250.200.4), and both devices have an assigned physical address (the router's is 00 00 A2 00 00 01 and host C's is 00 00 A2 00 10 40).

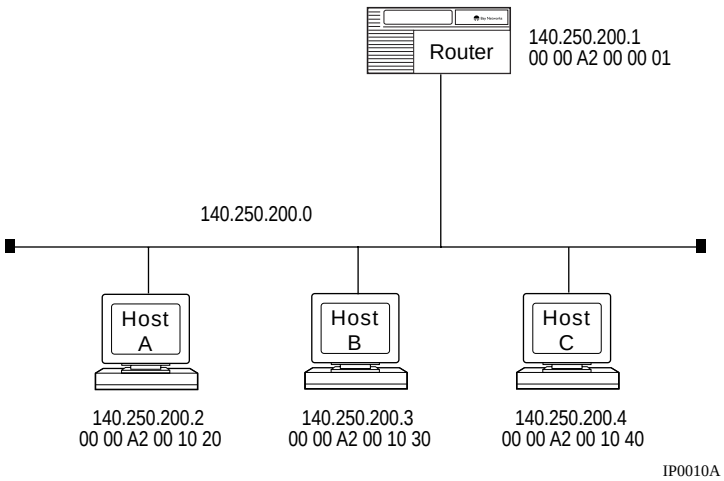

Figure 1-4. ARP Example

In Figure 1-4, the router wants to send a packet to host C but knows only host C's IP address. The router uses ARP to determine host C's physical address, as follows:

1. The router broadcasts a special packet, called an ARP request, that asks IP address 140.250.200.4 to respond with its physical address.
2. All network hosts receive the broadcast request.
3. Only host C responds with its hardware address.

The router maps host C's IP address (140.250.200.4) to its physical address (00 00 A2 00 10 40) and saves the results in an address-resolution cache for future use.

The router can send ARP requests even if ARP, which is a dynamically loaded module, is not currently loaded on the router. To ensure that ARP is loaded correctly on a slot using Site Manager, choose Events Manager > Options > Filters; then select LOADER and Debug, and choose File > Get Current Log File. Verify that ARP is loaded on a slot by locating the following message in the log:

```
#xx:01/01/99 10:10:55.00 DEBUG SLOT x LOADERCODE:33
Loader service completed for ARP.EXE 0xxxxxxxxx
```

Reverse Address Resolution Protocol (RARP)

You can configure a Nortel Networks router as a Reverse Address Resolution Protocol (RARP) server. A RARP server supplies clients on the same physical or logical LAN with IP addresses (Figure 1-5).

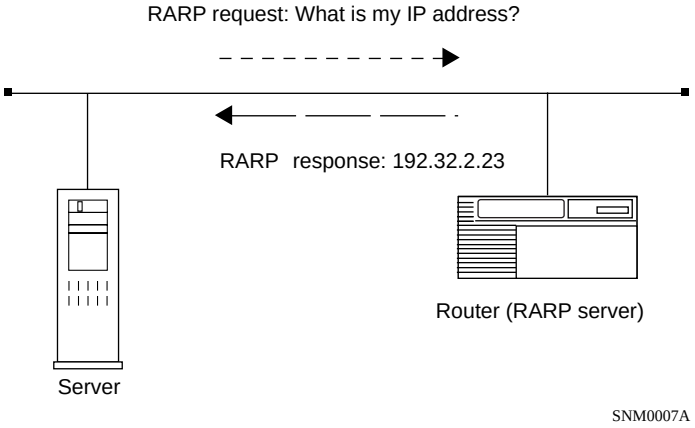

Figure 1-5. RARP Server Supplying an IP Address

To use RARP services, you must set up a MAC address-to-IP address mapping table. This table lists the MAC addresses of clients and the corresponding IP addresses that the RARP server assigns to those clients. When a client needs to acquire an IP address, the following interchange takes place:

1. The client broadcasts a RARP request specifying its MAC address.
2. Upon receiving a RARP request, the router refers to its MAC address-to-IP address mapping table, then sends the client a response packet containing the corresponding IP address.
3. The client examines the response packet to learn its IP address.

You can configure RARP support on Ethernet, token ring, and Fiber Distributed Data Interface (FDDI) interfaces.

Routing Information Protocol (RIP)

The Routing Information Protocol (RIP) is a distance-vector protocol that enables routers in the same autonomous system to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring networks and listen for RIP updates from the routers on those neighboring networks. Routers use the information in the RIP updates to keep their internal routing tables current. For RIP, the “best” path to a destination is the shortest path (the path with the fewest hops). RIP computes distance as a metric, usually the number of hops (or routers) from the source network to the target network.

Open Shortest Path First (OSPF) Protocol

The Open Shortest Path First (OSPF) protocol is an interior gateway protocol (IGP) intended for use in large networks. Using a link state algorithm, OSPF exchanges routing information between routers in an autonomous system. Routers synchronize their topological databases. After the routers are synchronized and the routing tables are built, the routers flood topology information only in response to some topological change. For OSPF, the “best” path to a destination is the path that offers the least cost metric delay. In OSPF, cost metrics are configurable, allowing you to specify preferred paths.

OSPF is a *link-state protocol*. A router running a link-state protocol periodically tests the status of the physical connection to each of its neighboring routers and sends this information to its other neighbors. A link-state protocol does not require each router to send its entire routing table to its neighbors. Instead, each OSPF router floods only link-state change information throughout the autonomous system (or area, if the AS is divided into areas). This process is referred to as the synchronization of the routers’ topological databases.

With the link-state information, each router builds a *shortest-path tree* with itself as the root of the tree. It can then identify the shortest path from itself to each destination and build its routing table.

OSPF supports CIDR and can carry supernet advertisements within a routing domain.

OSPF Addresses and Variable-Length Masks

A destination in an OSPF route advertisement is expressed as an IP address and a variable-length mask. Together, the address and the mask indicate the range of destinations to which the advertisement applies.

The ability to specify a range of networks allows OSPF to send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 128.185.0.0 with a mask of 0xffff0000 describes a single route to destinations 128.185.0.0 through 128.185.255.255.

OSPF Neighbors

OSPF *neighbors* are any two routers that have an interface to the same network. In each OSPF network, routers use the Hello protocol to discover their neighbors and maintain neighbor relationships. On a broadcast or point-to-point network, the Hello protocol dynamically discovers neighbors; however, on a nonbroadcast multi-access network, you must manually configure neighbors.

The Hello protocol ensures that communication between neighbors is bidirectional. Periodically, OSPF routers send out Hello packets over all interfaces. Included in these Hello packets is the following information:

- The router's priority
- The router's Hello timer and dead timer values
- A list of routers that have sent Hello packets to this router on this interface
- The router's choice for designated router and backup designated router

Bidirectional communication is verified when one router sees itself listed in the neighbor's Hello packet.

Neighbor Adjacencies

Neighbors may form an *adjacency* for the purpose of exchanging routing information. When two routers form an adjacency, they go through a process called *database exchange* to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent. From this point on, only routing change information is passed between the routers, thus conserving bandwidth.

Routers connected by a point-to-point network or a virtual link always form an adjacency. Also, every router on a multiaccess network forms an adjacency relationship with the designated router and the backup designated router.

Designated Routers

To further reduce the amount of routing traffic, the Hello protocol elects a *designated router* and a *backup designated router* on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link-state information with each other (which on a large network can mean a lot of routing protocol traffic), all routers on the network form adjacencies with the designated router and the backup designated router only and send link-state information to them. The designated router then distributes the information from each router to every other router.

The Hello protocol always elects a backup designated router along with the designated router. This router takes over all of the designated router's functions should the designated router fail.

OSPF Areas

OSPF routers reduce and restrict the amount of internal and external routing information that is flooded through the AS by dividing the AS into *areas*. Figure 1-6 shows an OSPF autonomous system divided into four areas and a required central area called a *backbone* that is used to distribute routing information among areas.

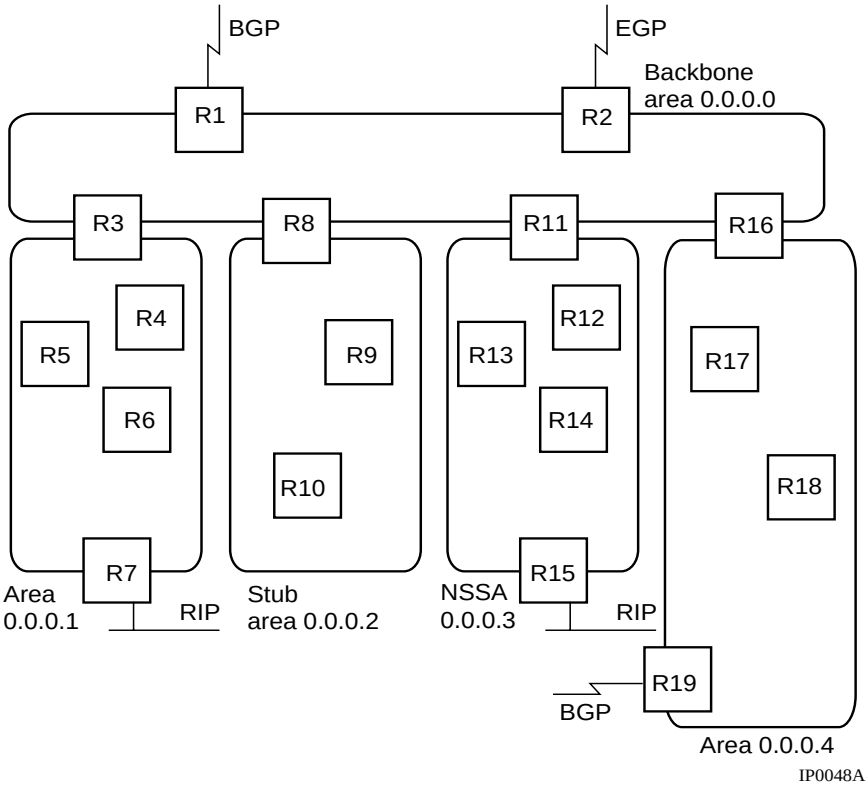

Figure 1-6. OSPF Areas

Each area has a unique ID number. (ID 0.0.0.0 is always reserved for the OSPF backbone.) The AS in Figure 1-6 has the following four areas:

- Backbone area: 0.0.0.0
- Normal areas: 0.0.0.1 and 0.0.0.4
- Stub area: 0.0.0.2
- Not-so-stubby area (NSSA): 0.0.0.3

For information about associating an OSPF interface with an area ID, see “Configuring an Area ID” on page 6-20.

Stub Area

A stub area does not import AS external routes and may or may not import internal route summaries. In place of routes to destinations outside the stub area, a border router connected to a stub injects a default route advertisement. When an internal router encounters a datagram addressed to a destination outside the stub area, the router forwards it to the border router specified in the default route advertisement.

Not-So-Stubby Area (NSSA)

A not-so-stubby area (NSSA) is similar to a stub area except that, to a limited degree, an NSSA can import AS external routes. Like a stub area, an NSSA consumes less memory and CPU resources by preventing the flooding of type 5 AS external link-state advertisements (LSAs) into the area and by relying on default routing to external destinations. However, unlike a stub area, an NSSA can import external routes into an OSPF routing domain. You can also configure type 7 address ranges to allow NSSA border routers to control the summarization of external routes for the NSSA.

OSPF Router Types

OSPF defines the following router types:

- *Internal router.* A router with interfaces to networks in one area only. Internal routers flood each area with complete routing information about changes that occur within the area. In Figure 1-6, R4, R5, and R6 in area 0.0.0.1 are internal routers.
- *Border router.* A router with an interface to the backbone network and interfaces to one or more additional areas. Each border router connects one or more areas to the backbone. In Figure 1-6, R3, R8, R11, and R16 are border routers. Using the backbone, border routers ensure that AS external (ASE) routes and summaries of routing information for all areas are distributed throughout the AS.
- *Boundary router.* A router configured to learn AS external routes and to inject this information into an OSPF AS. In Figure 1-6, R1, R2, R7, R15, and R19 are boundary routers.

AS External Routes

OSPF considers the following routes to be *AS external (ASE)* routes:

- A route to a destination outside the AS
- A static route
- A default route
- A route derived by RIP
- A directly connected network not running OSPF

In Figure 1-6, for example, routers R1 and R2 are boundary routers that use BGP and EGP to connect the backbone to external ASs. R7 in area 0.0.0.1 and R15 in area 0.0.0.3 are also boundary routers that use RIP to connect to external ASs. R19 in area 0.0.0.4 connects the area to an external AS via BGP.

OSPF Implementation Notes

This section provides suggestions to help you configure your OSPF network. The Nortel Networks implementation of OSPF does not restrict you to these suggestions, but we provide them as guidelines.

- Use the same password throughout an area, or even throughout the entire OSPF AS, if possible.
- Use the default timer values, unless you are running 9.6 KB synchronous lines. In this case, double the default timer values on both ends of the link.
- Use address ranges if your network is a subnetted network.
- Keep all subnets within one area. If you cross areas, you cannot configure summaries.
- Make sure the AS boundary router function is enabled if the router has any non-OSPF interfaces and if you want the router to propagate that information.
- Configure virtual links for each area border router that does not reside within or directly interface to the backbone. Every area border router must have a configured path to the backbone.
- OSPF considers the cost of a path, not just the hop count, when choosing the best path. Each interface, however, is assigned the default cost 1 for the path to which it interfaces. If you have a preferred path, you must change the metric cost for your interfaces. Assign a higher metric cost for those paths that are *not* preferred paths.
- If you have any devices in your network running OSPF, and you are now adding a Nortel Networks router, you must make sure that the router's timer values coincide with the timers in your other devices. Determine the timer values of the other devices, and change the router's timer values to match them.
- If you change the topology (for example, if you add an area, combine two areas, or move routers), you must reconfigure the appropriate OSPF elements (OSPF area ranges/interfaces/neighbors/virtual links, and so on).

Router Discovery Protocol

Before a host can send IP datagrams beyond its directly attached subnet, the host must discover the address of at least one operational router on that subnet. Router Discovery is an extension of the Internet Control Message Protocol (ICMP) that enables hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers.

Routers configured with Router Discovery periodically multicast or broadcast a route advertisement from each of their interfaces, announcing the IP address or addresses of that interface. Hosts discover the addresses of their neighboring routers by listening for these advertisements. Hosts will use the router with the highest preference level as a gateway.

Route Preferences

The IP router maintains an internal routing table. When determining how to forward a datagram, the IP router consults the table to determine the specific route that a datagram should take. A routing table can contain direct routes for the IP router's network interfaces, static routes, and any routes learned from RIP, OSPF, BGP, or EGP. (Information about adjacent hosts is maintained in a separate table.)

When a routing table contains multiple routes to the same destination, IP uses (among other information) a preference value to determine which route to select. Preference values range from 1 through 16 (the higher the number, the greater the preference).

By default, RIP, BGP, EGP, and OSPF external routes have a preference value of 1. Static routes, direct routes, and OSPF intra-area and interarea routes have a default preference value of 16.

You can configure a preference value from 1 through 16 for RIP, BGP, EGP, OSFP external, and static routes. You cannot configure the preference value of direct routes and OSPF intra-area and interarea routes.

To assign a preference value to a route learned by RIP, OSPF, BGP, or EGP, you configure an accept policy for the route (see "IP Routing Policies and Filters" on page 1-22). If an incoming route matches the policy, IP assigns the preference value you specify to the route and considers the route for possible inclusion in the routing table.

Route Weights

Route-weight calculation is an internal tool that IP uses to facilitate the selection of the best route among alternative routes to the same destination. Route-selection criteria are encoded into the route weight in a way that allows IP to compare routes simply by comparing their weight values, regardless of route sources.

Appendix D contains a worksheet that you can use to calculate route weights in your configuration.

Route-weight calculation increases the efficiency of the route-selection process. It also reduces the size of the routing database because all route selection parameters for each route are encoded in a single integer -- the weight value -- rather than stored in separate variables.

Using selection criteria encoded in the route weight, IP chooses routes in the following order:

1. The route with the highest preference value (see “Route Preferences” on page 1-20)
2. A direct route with the lowest metric
3. An OSPF intra-area route with the lowest metric
4. An OSPF interarea route with the lowest metric
5. An OSPF type 1 external route with the lowest metric
6. A BGP route with the highest LOCAL_PREF value
7. A RIP route with the lowest metric
8. An EGP route with the lowest metric
9. A static route with the lowest metric

Note: If OSPF is configured to propagate external routes using the route weight as the type 2 metric, routes that are received as OSPF ASE type 2 routes are evaluated according to their respective origins (for example, RIP or BGP).

IP Routing Policies and Filters

The IP router allows you to control the flow of routing data to and from the routing tables. This control is provided by two mechanisms:

- IP accept and announce policies
- IP import and export route filters

Note: Accept and announce policies provide a superset of the parameters provided by import and export route filters. Nortel Networks supports both IP policies and IP route filters. However, if you are using import and export filters for routing table management, you should migrate as soon as possible to IP policies.

IP *accept policies* (and the subset of parameters provided by import filters) govern the addition of new routes (derived from RIP, OSPF, BGP, or EGP) to the routing tables. When RIP, OSPF, BGP, or EGP receives a new routing update, it consults its accept policies to validate the information before entering the update into the routing tables. Accept policies contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matching routes).

IP *announce policies* (and the subset of parameters provided by export filters) govern the propagation of RIP, OSPF, BGP, or EGP routing information. When preparing a routing advertisement, RIP, OSPF, BGP, or EGP consults its announce policies to determine whether the routes to specific networks are to be advertised and how they are to be propagated. Announce policies contain network numbers (to associate a policy with a specific network) and action information (to specify a route propagation procedure).

For information about BGP and EGP, see *Configuring IP Exterior Gateway Protocols (BGP and EGP).*

Every IP router maintains a table of current routing information. The routing table manager receives routing updates from the network through the IP routing protocols running on the router. Periodically, the routing table manager issues routing updates through the protocols.

Figure 1-7 shows a router configured with all of the IP routing protocols supported by Nortel Networks: OSPF, RIP, BGP-3, BGP-4, and EGP. The arrows indicate the direction of flow of routing information between the network and the protocols running on the router, between the protocols and the routing table manager, and between the routing table manager and the routing table.

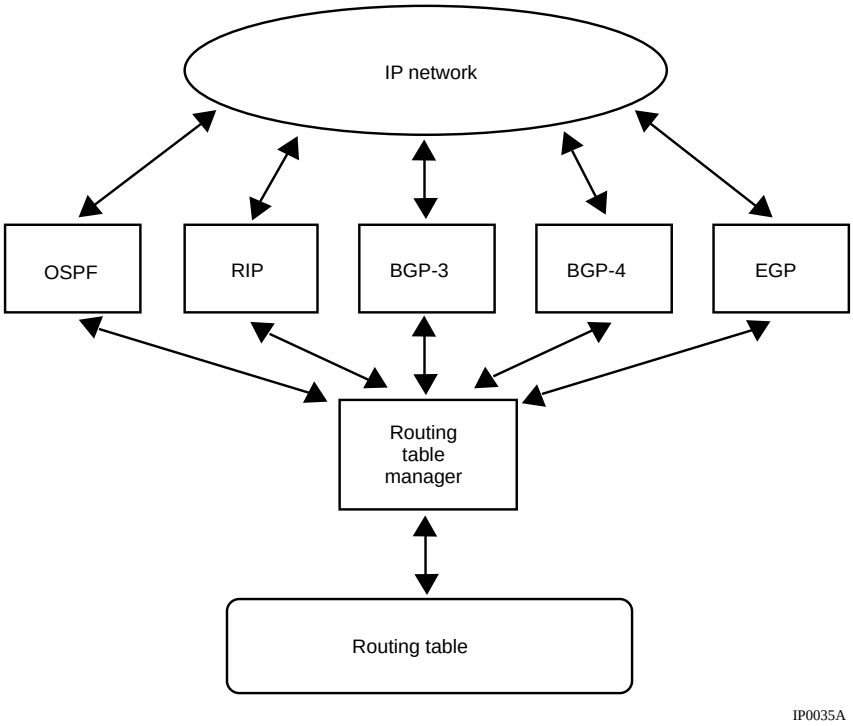

Figure 1-7. IP Routing Table

The flow of routing information between the network, the protocols, and the routing table manager is controlled by *routing information policies.*

Each time that a routing update arrives from a remote router, the following steps occur (see Figure 1-8):

1. The protocol receiving the route consults an *accept policy* to determine whether to forward the route to the IP routing table manager or drop the route.
2. If the protocol forwards the route, the routing table manager determines whether to inject the route into the routing table.

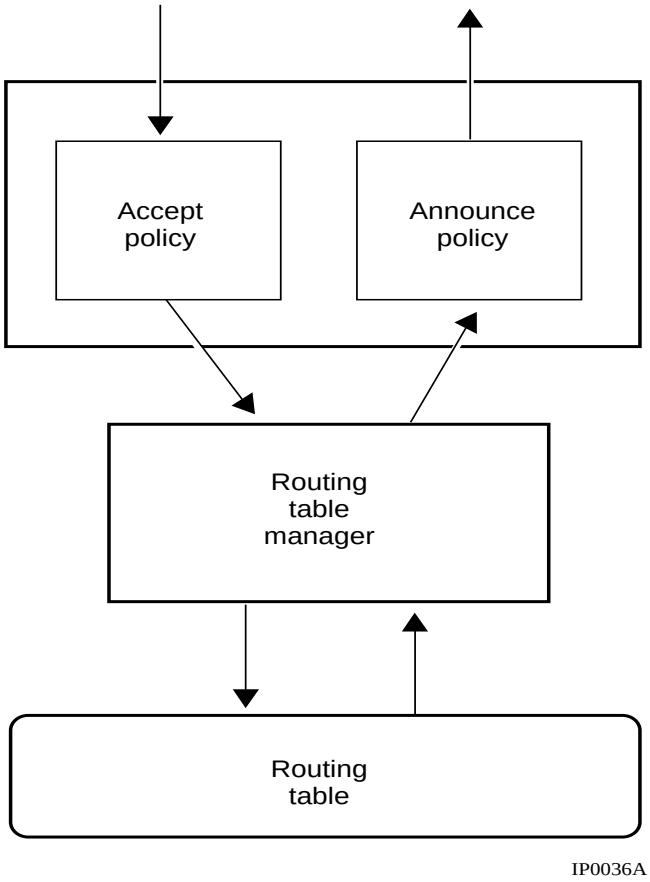

Figure 1-8. Accept and Announce Policies

Periodically, the routing table manager announces routes to other routers in the network as follows:

1. The routing table manager forwards a route for advertisement to the protocol.
2. The protocol consults an *announce policy* to determine whether or not to advertise the route to the network.

For instructions on configuring RIP policies, see “Configuring RIP Accept and Announce Policies” on page 5-24.

For descriptions of IP import and export filter parameters, see Appendix C.

How OSPF Implements Routing Policies

The way OSPF applies accept and announce policies to routing information differs in several ways from the procedure shown in Figure 1-7. OSPF link state advertisements (LSAs) are received and placed in the link state database (LSDB) of the router. The information in the LSDB is also propagated to other routers in the OSPF routing domain.

According to the OSPF standard, all routers in a given area must maintain a similar database. To maintain database integrity across the network, a router must not manipulate received LSAs before propagating them to other routers. To accomplish this, OSPF accept and announce policies act in the following manner:

- OSPF accept policies control which OSPF non-self-originated external routing information is passed to the routing table manager. The accept policies control only what the local router uses; they do not affect the propagation of OSPF internal and OSPF non-self-originated external information to other routers.
- OSPF announce policies control which self-originated external routing updates are placed into the LSDB for distribution according to the OSPF standard. OSPF announce policies affect what other routers learn but only with regard to the local router’s self-originated information.

For instructions on configuring OSPF policies, see “Configuring OSPF Accept and Announce Policies” on page 6-57.

IP Traffic Filters

A traffic filter enables the router to selectively relay or drop an inbound packet, frame, or datagram based on standard protocol fields or user-defined fields. Traffic filters apply to incoming traffic only.

For information about IP traffic filters, see *Configuring Traffic Filters and Protocol Prioritization.*

Chapter 2
Starting IP Services

This chapter describes how to use the BCC and Site Manager to perform a basic configuration -- that is, a configuration using all available defaults -- for the IP services described in this guide.

Topic	Page
Using the BCC	2-2
Using Site Manager	2-5

Using the BCC

You can use the BCC to start IP, RIP, OSPF, and Router Discovery. The following sections provide instructions for starting these services.

Topic	Page
Starting IP	2-2
Starting RIP	2-3
Starting OSPF	2-4
Starting Router Discovery	2-4

Starting IP

To start IP on the router, you must:

1. Configure a physical interface on an available slot/connector.
2. Configure an IP interface on the physical interface.

Step 1: Configuring a Physical Interface

To configure a physical interface on a slot and connector, navigate to the top-level box prompt (box or stack) and enter:

<interface_type> **slot** *<slot_number>* [**module** *<module_number>*] **connector** *<connector_number>*

interface_type is the name of a link module on the router.

slot_number is the number of the slot on which the link or net module is located.

module_number is the net module number (ASN™ and System 5000™ routers only).

connector_number is the number of a connector on the link or net module.

For example, the following command configures an Ethernet interface on slot 2, connector 2:

```
box# ethernet slot 2 connector 2
ethernet/2/2#
```

Step 2: Configuring an IP Interface

To configure an IP interface on a physical interface, navigate to the prompt for the physical interface and enter:

ip address <*address*> **mask** <*mask*>

address and *mask* are a valid IP address and its associated mask, expressed in dotted-decimal notation.

For example, the following command configures IP interface 2.2.2.2/255.0.0.0 on an Ethernet physical interface on slot 2, connector 2:

```
ethernet/2/2# ip address 2.2.2.2 mask 255.0.0.0
ip/2.2.2.2/255.0.0.0#
```

An IP interface is now configured on the Ethernet interface with default values for all interface parameters. When you configure an IP interface, the BCC also configures IP globally on the router with default values for all IP global parameters.

You customize IP by modifying IP global and interface parameters as described in Chapter 3, “Configuring and Customizing IP.”

Starting RIP

You start RIP on the router by adding RIP to an existing IP interface.

To start RIP, navigate to an IP interface prompt and enter:

rip

For example, the following command adds RIP to IP interface 2.2.2.2/255.0.0.0:

```
ip/2.2.2.2/255.0.0.0# rip
rip/2.2.2.2#
```

RIP is now running on the router and on the interface with default values for all parameters. You customize RIP by modifying RIP parameters as described in Chapter 5, “Customizing RIP Services.”

Starting OSPF

You start OSPF on the router by adding OSPF to an existing IP interface.

To start OSPF, navigate to an IP interface prompt and enter:

ospf area *<area_id>*

area_id is the identifier, expressed in dotted-decimal notation, of the OSPF area to which the router is connected through the IP interface.

For example, the following command adds OSPF to IP interface 2.3.3.3/255.0.0.0. This interface connects the router to OSPF area 0.0.0.0, the OSPF backbone.

```
ip/2.3.3.3/255.0.0.0# ospf area 0.0.0.0
ospf/2.3.3.3#
```

OSPF is now running on the router and on the interface with default values for all parameters. You customize OSPF by modifying OSPF parameters as described in Chapter 6, “Customizing OSPF Services.”

Starting Router Discovery

You start Router Discovery by adding it to an IP interface.

To start Router Discovery, navigate to an IP interface prompt and enter:

rdisc

For example, the following command adds Router Discovery to IP interface 2.2.2.2/255.0.0.0:

```
ip/2.2.2.2/255.0.0.0# rdisc
rdisc/2.2.2.2#
```

Router Discovery is now running on IP interface 2.2.2.2/255.0.0.0 with default values for all parameters. You customize Router Discovery by modifying parameters as described in “Configuring and Customizing Router Discovery” on page 3-54.

Using Site Manager

You can use Site Manager to start IP, RIP, OSPF, and RARP, and to configure a circuitless IP interface and an unnumbered interface. The following sections provide instructions.

Topic	Page
Starting IP	2-5
Starting RIP	2-6
Starting OSPF	2-8
Starting RARP	2-9
Configuring a Circuitless IP Interface	2-10
Configuring an Unnumbered IP Interface	2-11

Starting IP

Before you can select a protocol to run on the router, you must first configure a circuit that the protocol can use as an interface to an attached network. For information and instructions, see *Configuring WAN Line Services* and *Configuring Ethernet, FDDI, and Token Ring Services.*

When you have successfully configured the circuit, the Select Protocols window opens. To start IP, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Select Protocols window, select **IP**. Then click on **OK**.	The IP Configuration window opens.
2. Set the following parameters: • **IP Address** • **Subnet Mask** • **Transmit Bcast Addr** • **UnNumbered Assoc Address** Click on **Help** or see the parameter descriptions beginning on page A-2.	
3. Click on **OK**.	You return to the Configuration Manager window.

An IP interface is now configured with default values for all interface parameters. When you configure an IP interface, Site Manager also configures IP globally on the router with default values for all IP global parameters.

You customize IP by modifying IP global and interface parameters as described in Chapter 3, “Configuring and Customizing IP.”

For information about configuring unnumbered interfaces, see “Configuring an Unnumbered IP Interface” on page 2-11.

Starting RIP

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as interface to an attached network. For information and instructions, see *Configuring WAN Line Services* and *Configuring Ethernet, FDDI, and Token Ring Services.*

When you have successfully configured the circuit, the Select Protocols window opens. Complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Select Protocols window, select the following protocols: • **IP** • **RIP** Then click on **OK**.	The IP Configuration window opens.
2. Set the following parameters: • **IP Address** • **Subnet Mask** • **Transmit Bcast Addr** • **UnNumbered Assoc Address** Click on **Help** or see the parameter descriptions beginning on page A-2.	
3. Click on **OK**.	You return to the Configuration Manager window.

RIP is now running on the router and on the interface with default values for all parameters. You customize RIP by modifying RIP parameters as described in Chapter 5, “Customizing RIP Services.”

Adding RIP to an IP Interface

To add RIP to an existing IP interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on the connector to which you want to add RIP services.	The Edit Connector window opens.
2. Click on **Edit Circuit**.	The Circuit Definition window opens.
3. Choose **Protocols**.	The Protocols menu opens.
4. Choose **Add/Delete**.	The Select Protocols window opens.
5. Click on **RIP**.	Site Manager highlights the selection.
6. Click on **OK**.	You return to the Circuit Definition window.
7. Choose **File**.	The File menu opens.
8. Choose **Exit**.	You return to the Configuration Manager window.

Starting OSPF

Before you can select a protocol to run on the router, you must configure a circuit that the protocol can use as interface to an attached network. For information and instructions, see *Configuring WAN Line Services* and *Configuring Ethernet, FDDI, and Token Ring Services.*

After you have successfully configured the circuit, the Select Protocols window opens. To start OSPF, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Select Protocols window, select the following protocols: • **IP** • **OSPF/MOSPF** Then click on **OK**.	The IP Configuration window opens.
2. Set the following parameters: • **IP Address** • **Subnet Mask** • **Transmit Bcast Addr** • **UnNumbered Assoc Address** Click on **Help** or see the parameter descriptions beginning on page A-2.	
3. Click on **OK**.	The Initial OSPF Global Configuration window opens.
4. Click on **OK** to accept the default values.	The OSPF Area Address window opens.
5. Set the **Area Address** parameter. Click on **Help** or see the parameter description on page A-43.	
6. Click on **OK**.	The Broadcast Type window opens.
7. Set the **Broadcast Type** parameter. Click on **Help** or see the parameter description on page A-43.	
8. Click on **OK**.	You return to the Configuration Manager window.

OSPF is now running on the router and on the interface with default values for all parameters. You customize OSPF by modifying OSPF parameters as described in Chapter 6, “Customizing OSPF Services.”

Starting RARP

Before you can enable RARP services, you must enable IP on the router. You can, however, enable IP and RARP services on the router simultaneously. If you want to add RARP to a circuit on which you have already configured IP, see *Configuring and Managing Routers with Site Manager* for information about adding a protocol.

To enable RARP on an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, select the link- or net-module connector on which you are enabling RARP services.	The Select Protocols window opens.
2. Select **Reverse ARP**. When you select Reverse ARP, you automatically select IP.	
3. Click on **OK**.	The IP Configuration window opens.
4. Set the following parameters: • **IP Address** • **Subnet Mask** • **Transmit Bcast Addr** • **UnNumbered Assoc Address** Click on **Help** or see the parameter descriptions beginning on page A-2.	
5. Click on **OK**.	You return to the Configuration Manager window.

RARP is now running on the router and on the interface with default values for all parameters. You customize RARP by modifying RARP parameters as described in “Customizing RARP” on page 4-12.

Configuring a Circuitless IP Interface

A *circuitless* IP interface has an IP address that is not mapped to a specific circuit. If one or more of the router's IP interfaces become disabled, the circuitless IP interface address ensures that the router is always reachable as long as a viable path to the router exists. The IP router can support one circuitless IP interface.

IP traffic is delivered to and transmitted from the circuitless interface in the same way as any other IP interface. In addition, the circuitless IP interface can receive packets from any application.

When you configure a circuitless IP interface, note the following:

- You can configure one circuitless IP interface per router. Additional circuitless IP interfaces will not initialize.
- You can add OSPF and BGP to a circuitless interface.
- You must assign a unique IP address and subnetwork number to the circuitless IP interface.
- You cannot configure a circuitless IP interface in nonforwarding mode.

To configure a circuitless IP interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Circuitless IP**.	The Circuitless IP menu opens.
4. Choose **Create**.	The IP Configuration window opens.
5. Click on **OK**.	Site Manager saves the circuitless IP interface and opens a special Select Protocols window that lists the protocols that you can configure on a circuitless interface.
6. Choose one or more protocols and click on **OK**.	A configuration window opens for each protocol that you selected.
7. Configure the protocol by setting the parameters. Click on **Help** or see the protocol-specific guide.	

Choosing Slots to Support the Circuitless Interface

By default, all slots support the IP circuitless interface. To specify the slots that can support the circuitless interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the circuitless interface.	The window displays the parameter values for that interface.
5. Set the **Slot Mask** parameter. Click on Help or see the parameter description on page A-15.	
6. Click on **Apply**, and then click on **Done.**	You return to the Configuration Manager window.

Configuring an Unnumbered IP Interface

IP allows you to configure an interface on a point-to-point connection without using an IP address. Such an interface is called an *unnumbered interface*. The ability to establish a point-to-point link using an unnumbered IP interface helps alleviate two of the major problems caused by the continued rapid growth of the Internet: exhaustion of Class B network addresses and of the 32-bit IP address space.

You associate each unnumbered interface with the IP address of any numbered interface on the router, including the circuitless interface. The router can support multiple unnumbered interfaces. Multiple unnumbered interfaces can be associated with the same IP address. Because all traffic over an unnumbered interface uses broadcast addressing at the link layer, neither an adjacent host specification nor address resolution is required.

As it does with routes learned over numbered interfaces, IP stores each route learned over an unnumbered interface in the routing table. The routing-table entry for a route learned over an unnumbered interface contains the following values:

Next-hop address	0
Next-hop mask	0
Next-hop interface	Circuit number of the unnumbered interface

Note: Unnumbered interfaces cannot be pinged directly. For this reason, such interfaces can make it difficult to diagnose router problems.

Point-to-point connections using unnumbered interfaces can be configured to advertise RIP, OSPF, IBGP, DVMRP, and static routes.

The associated address assigned to the unnumbered interface determines whether or not RIP configured to send updates in Version 1 mode will advertise a subnetwork over the unnumbered interface. The associated address also determines which mask is applied to RIP V1 updates received on that interface. For unnumbered links using RIP V1, the defined associated addresses at each end of the link must belong to the same network and have the same mask for routes to be exchanged correctly.

If a subnetwork on the router has the same mask as the associated address, RIP V1 will advertise that subnet over the unnumbered interface. If the mask on the subnetwork is different from the mask of the associated address, RIP V1 advertises only the natural network of the subnet.

Note: Nortel Networks recommends that you select RIP2 mode for unnumbered interfaces. With RIP2, RIP updates contain both the route and mask information.

BGP peers, NetBIOS, and BootP cannot be configured directly on an unnumbered interface. To configure a BGP peer-to-peer session on routers connected through unnumbered interfaces, see *Configuring IP Exterior Gateway Protocols (BGP and EGP).* To route NetBIOS packets over an unnumbered interface, you must configure a static entry to the name server. To run BootP over unnumbered interfaces, you must select a preferred BootP server. For instructions, see *Configuring SNMP, BootP, and DHCP Services.*

Using the Alternate Associated Address Option

The alternate associated address option ensures that a network on an unnumbered interface remains reachable. IP automatically assigns an alternate associated address to an unnumbered interface in the event that the primary associated address goes down. IP uses the first available interface.

Note: If an associated address becomes unreachable, some functionality may be lost for certain protocols over the unnumbered interface.

To enable the associated address option, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	The window displays the parameter values for that interface.
5. Set the **Unnumbered Associated Alternate** parameter. Click on Help or see the parameter description on page A-17.	
6. Click on **Apply**, and then click on **Done.**	You return to the Configuration Manager window.

Note: In some configurations, changing the associated address can affect the way routes are advertised. For example, if you change the associated address for an unnumbered interface configured with RIP, you may change the way RIP advertises subnets.

Chapter 3
Configuring and Customizing IP

You customize IP services by setting parameters as described under the following topics:

Topic	Page
Customizing IP Global Parameters	3-2
Customizing an IP Interface	3-20
Configuring an Adjacent Host Address	3-46
Defining a Static Route	3-49
Configuring and Customizing Router Discovery	3-54

Customizing IP Global Parameters

When you configure an IP interface on a slot, IP is automatically configured globally on the slot with default values for all global parameters. You customize global IP by modifying global IP parameters as described under the following topics:

Topic	Page
Disabling and Reenabling Global IP	3-3
Configuring the Router for Not-Forwarding Mode	3-4
Configuring Bridging on a Router in Not-Forwarding Mode	3-6
Setting the Time-to-Live Value on a Source Packet	3-9
Allowing an All-Zero or All-One Subnet Address	3-10
Estimating the Size of the Routing Table	3-11
Using a Default Route for an Unknown Subnet	3-12
Specifying the Maximum Number of IP Policies	3-13
Disabling and Reenabling Route Filter Support	3-14
Configuring Equal-Cost Multipath Support for RIP and OSPF	3-15
Customizing the IP Routing Table Structure	3-18
Specifying the Percentage of Buffers Available to ARP	3-19

Disabling and Reenabling Global IP

IP is enabled on the slot by default. You can disable and reenable IP as required.

Using the BCC

To disable or reenable IP on the slot, navigate to the IP global prompt (for example, **box; ip**) and enter:

state <*state*>

state is one of the following:

enabled (default)
disabled

For example, the following command disables IP on the router:

```
ip# state disabled
ip#
```

Using Site Manager

To disable or reenable IP on the slot, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-19.	
5. Click on **OK**.	You return to the Configuration Manager window.

Configuring the Router for Not-Forwarding Mode

By default, IP forwards all packets that are not addressed to itself. You can also configure IP in *not-forwarding* -- or *host-only* -- mode.

Use the forwarding mode if you want the IP router to route (forward) IP traffic. In forwarding mode, the IP router processes all broadcast packets and all IP packets explicitly addressed to it, and routes all other IP packets.

Choose not-forwarding mode on the router if you want to provide IP management access (by means of TFTP and SNMP) to all active IP interfaces but also want to prohibit the IP router from forwarding IP traffic. You must specify an identical IP address and mask combination for each active IP interface that will provide management access.

In not-forwarding mode, the IP router acts as an IP host; it does not forward IP traffic, but it still processes packets explicitly addressed to it. In not-forwarding mode, only static routes and adjacent-host routes are allowed. No routing protocols are initiated.

You can use the BCC or Site Manager to select the forwarding mode.

Using the BCC

To specify the forwarding mode, navigate to the IP global prompt (for example, **stack; ip**) and enter:

forwarding <*mode*>

mode is one of the following:

forwarding (default)
notforwarding

For example, the following command puts the router in not-forwarding mode:

```
ip# forwarding notforwarding
ip#
```

Using Site Manager

To specify the forwarding mode, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Forwarding** parameter. Click on **Help** or see the parameter description on page A-19.	
5. Click on **OK**.	You return to the Configuration Manager window.

Configuring Bridging on a Router in Not-Forwarding Mode

Because the IP router does not forward IP traffic in not-forwarding mode, you must configure the router to bridge IP traffic not explicitly addressed to it. You must configure the bridge for each circuit that conveys IP datagrams. The bridge will then forward all IP datagrams that are not explicitly addressed to the router.

Table 3-1 and Table 3-2 show valid and invalid configurations for source routing bridges and learning bridges. Each configuration has the following format:

source_device > medium > destination_device

or

source_device > medium1 > intermediate_device > medium2 > destination_device

Table 3-1. Source Routing Bridge Support for Host-Only Mode

Bridge Configuration	Support
Bay* > Eth† > Bay	Supported
Bay > Token‡ > Bay	Supported
Bay > FDDI** > Bay	Supported
Bay > PTP†† > Bay	Supported
Bay > FR‡‡ > Bay	Supported
Bay > SMDS*** > Bay	Supported
Bay > PPP††† > Bay	Not supported
Bay > Eth > Bay >Token > ES‡‡‡	Not supported
Bay > Token > Bay > Token > ES	Supported
Bay > FDDI > Bay > Token > ES	Not supported
Bay > PTP > Bay > Token > ES	Not supported
Bay > FR > Bay > Token > ES	Not supported
Bay > SMDS > Bay > Token > ES	Not supported
Bay > PPP > Bay > Token > ES	Not supported
ES > Token > Bay > Eth > Bay	Not supported
ES > Token > Bay > Token > Bay	Supported

(continued)

Table 3-1. Source Routing Bridge Support for Host-Only Mode *(continued)*

Bridge Configuration	Support
ES > Token > Bay > FDDI > Bay	Not supported
ES > Token > Bay > PTP > Bay	Not supported
ES > Token > Bay > FR > Bay	Not supported
ES > Token > Bay > SMDS > Bay	Not supported
ES > Token > Bay > PPP > Bay	Not supported
ES > Token > Bay > Eth > Bay > Token > ES	Supported
ES > Token > Bay > Token > Bay > Token > ES	Supported
ES > Token > Bay > FDDI > Bay > Token > ES	Supported
ES > Token > Bay > PTP > Bay > Token > ES	Supported
ES > Token > Bay > FR > Bay > Token > ES	Supported
ES > Token > Bay > SMDS > Bay > Token > ES	Supported
ES > Token > Bay > PPP > Bay > Token > ES	Supported

* Nortel Networks router with bridge and IP in host-only mode

† Ethernet connection

‡ Token ring connection

** FDDI connection

†† Nortel Networks proprietary point-to-point synchronous connection

‡‡ Frame relay synchronous connection

*** SMDS synchronous connection

††† PPP synchronous connection

‡‡‡ Station you are communicating with if not Nortel Networks

Table 3-2. Learning Bridge Support for Host-Only Mode

Bridge Configuration	Support
Bay* > Eth† > Bay	Supported
Bay > Token‡ > Bay	Supported
Bay > FDDI** > Bay	Supported
Bay > PTP†† > Bay	Supported
Bay > FR‡‡ > Bay	Supported
Bay > SMDS*** > Bay	Supported

(continued)

Table 3-2. Learning Bridge Support for Host-Only Mode *(continued)*

Bridge Configuration	Support
Bay > PPP††† > Bay	Not supported
Bay > Eth > Bay > Eth > ES‡‡‡	Supported
Bay > Token > Bay > Eth > ES	Not supported
Bay > FDDI > Bay > Eth > ES	Supported
Bay > PTP > Bay > Eth > ES	Supported
Bay > FR > Bay > Eth > ES	Not supported
Bay > SMDS > Bay > Eth > ES	Not supported
Bay > PPP > Bay > Eth > ES	Not supported
ES > Eth > Bay > Eth > Bay	Supported
ES > Eth > Bay > Token > Bay	Not supported
ES > Eth > Bay > FDDI > Bay	Supported
ES > Eth > Bay > PTP > Bay	Supported
ES > Eth > Bay > FR > Bay	Not supported
ES > Eth > Bay > SMDS > Bay	Not supported
ES > Eth > Bay > PPP > Bay	Not supported
ES > Eth > Bay > Eth > Bay > Eth > ES	Supported
ES > Eth > Bay > Token > Bay > Eth > ES	Supported
ES > Eth > Bay > FDDI > Bay > Eth > ES	Supported
ES > Eth > Bay > PTP > Bay > Eth > ES	Supported
ES > Eth > Bay > FR > Bay > Eth > ES	Supported
ES > Eth > Bay > SMDS > Bay > Eth > ES	Supported
ES > Eth > Bay > PPP > Bay > Eth > ES	Supported

* Nortel Networks router with bridge and IP in host-only mode

† Ethernet connection

‡ Token ring connection

** FDDI connection

†† Nortel Networks proprietary point-to-point synchronous connection

‡‡ Frame relay synchronous connection

*** SMDS synchronous connection

††† PPP synchronous connection

‡‡‡ Station you are communicating with if not Nortel Networks

Setting the Time-to-Live Value on a Source Packet

Each IP data packet includes a *time-to-live* (TTL) value. The TTL value specifies the maximum number of hops that the packet is allowed to traverse in the network before an intermediate router discards the packet.

The router that originates the packet sets the TTL to a positive value. Each router that receives the packet decrements the TTL counter by one. A router that receives a packet with a TTL of zero discards the packet if the packet is not addressed to itself. The TTL counter prevents packets from looping endlessly through the network.

By default, IP sets the TTL field on each source packet (that is, each packet that it originates and transmits) to 30 hops. You can use the BCC or Site Manager to set the TTL value as required.

Using the BCC

To set the TTL value, navigate to the global IP prompt (for example, **box; ip**) and enter:

ttl *<max_hops>*

max_hops is the maximum number of hops that the packet can traverse before an intermediate router discards it.

For example, the following command sets the TTL value to 25 hops:

```
ip# ttl 25
ip#
```

Using Site Manager

To set the TTL value, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Default TTL** parameter. Click on **Help** or see the parameter description on page A-21.	
5. Click on **OK**.	You return to the Configuration Manager window.

Allowing an All-Zero or All-One Subnet Address

By default, an IP address with a subnet portion of all zeros or all ones is an illegal address.

You can configure IP to allow an all-zero or all-one subnet address. Enable this feature with caution, however, for it can result in an ambiguous address. For example, if an all-zero subnet address and an all-zero broadcast address are both valid, the router cannot distinguish an all-subnets broadcast from a directed broadcast for the zero subnet.

Using the BCC

To allow an all-zero or all-one subnet address, navigate to the IP global prompt (for example, **box; ip**) and enter:

all-subnets enabled

For example:

```
ip# all-subnets enabled
ip#
```

Using Site Manager

To allow an all-zero or all-one subnet address, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Zero Subnet Enable** parameter. Click on **Help** or see the parameter description on page A-22.	
5. Click on **OK**.	You return to the Configuration Manager window.

Estimating the Size of the Routing Table

IP allows you to estimate how many networks and hosts require an entry in the IP routing table. The router uses your estimate to preallocate memory for the routing table. Preallocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation. Preallocation also makes better use of memory and reduces the amount of memory required.

By default, the router allocates resources to support 500 network and host entries in the routing table. Avoid making an estimate that is excessively large; doing so will cause a wasteful overallocation of memory.

To preallocate memory for the routing tables, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the following parameters: • **Estimated Networks** • **Estimated Hosts** Click on **Help** or see the parameter descriptions beginning on on page A-23.	
5. Click on **OK**.	You return to the Configuration Manager window.

Using a Default Route for an Unknown Subnet

By default, IP drops packets addressed to an unknown subnet and returns an ICMP message to the sender. This prevents local traffic from accidentally following the default route to the Internet.

In cases where remote sites follow a default route to a central site, it is appropriate to enable this parameter. If the router serves as an Internet gateway (with a default route to the Internet), the parameter can be disabled.

The default route must be present in the routing table.

To allow the IP router to use a default route for unknown subnets, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Enable Default Route for Subnets** parameter. Click on **Help** or see the parameter description on page A-24.	
5. Click on **OK**.	You return to the Configuration Manager window.

Specifying the Maximum Number of IP Policies

By default, IP allows you to configure up to 32 announce policies and 32 accept policies for each protocol that you configure on the router. You must increase this value if you want IP to allocate more memory and implement additional policies.

Using the BCC

To set the maximum number of policies, navigate to the IP global prompt (for example, **stack; ip**) and enter:

max-policies *<max_policies>*

max_policies is the maximum number of accept or announce policies that you can configure for each routing protocol.

For example, the following command sets the maximum number of policies to 50:

```
ip# max-policies 50
ip#
```

Using Site Manager

To set the maximum number of policies, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Maximum Policy Rules** parameter. Click on **Help** or see the parameter description on page A-24.	
5. Click on **OK**.	You return to the Configuration Manager window.

Disabling and Reenabling Route Filter Support

By default, IP supports route filters. When route filter support is disabled, IP does not allocate memory for route filters when the maximum number of IP policies is increased. You can use the BCC or Site Manager to disable and reenable this feature as required.

Using the BCC

To disable or reenable support for route filters, navigate to the IP global prompt (for example, **box; ip**) and enter:

route-filters *<state>*

state is one of the following:

enabled (default)
disabled

Using Site Manager

To disable or reenable support for route filters, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Route Filter Support** parameter. Click on **Help** or see the parameter description on page A-24.	
5. Click on **OK**.	You return to the Configuration Manager window.

Configuring Equal-Cost Multipath Support for RIP and OSPF

By default, the IP routing table stores a single best RIP route and a single best OSPF route to a given destination. If either protocol submits another route to the same destination, IP compares the new route with the current route. If the new route is better, IP replaces the current route with the new route. If not, IP discards the new route.

IP *equal-cost multipath (ECMP)* support is a load-balancing feature that allows IP to distribute traffic over multiple (up to five) equal-cost paths to the same destination.

IP supports three methods of distribution for equal-cost routes:

- *Round-robin* distribution. IP forwards each packet to a different next hop until it reaches the end of the list of available next hops; then it repeats the list. Round-robin distribution makes full use of available resources but may cause packets to be delivered out of order.
- *Source-destination hash* distribution based on the source and destination address. IP forwards all packets with a given source and destination address to the same next hop. This method increases the chances that the packets will be delivered in order.

- *Destination-hash* distribution based on the destination address only. IP forwards all packets with a given destination address to the same next hop.

If you enable equal-cost multipath support on the router, IP can store multiple equal-cost best RIP and OSPF routes in the routing table. When RIP or OSPF submits a route to a destination, one of the following events occurs:

- IP determines that the current route to that destination is better than the new route. IP discards the new route.
- IP determines that the new route is better than the current route. IP discards the current route and replaces it with the new route. If the routing table contains multiple equal-cost best routes, IP discards all of these routes.
- IP determines that the new route and the current route have the same cost. IP adds the new route to the routing table -- up to a maximum number that you specify. If the routing table already contains the maximum number of equal-cost routes from RIP or OSPF, IP discards the route.

By default, equal-cost multipath support is disabled on the router. If you enable equal-cost multipath support on the router, you must set the maximum number of equal-cost routes (up to five) that IP stores in the routing table for RIP and OSPF.

Using the BCC

To enable or disable equal-cost multipath support, navigate to the IP global prompt (for example, **box; ip**) and enter:

ecmp-method *<method>*

method is one of the following:

disabled (default)
roundRobin
srcDestHash
destinationHash

To specify the maximum number of equal-cost paths for RIP, navigate to the IP global prompt (for example, **box; ip**) and enter:

rip-max-paths *<max_number>*

max_number is an integer from 1 (the default) through 5.

To specify the maximum number of equal-cost paths for OSPF, navigate to the IP global prompt (for example, **box; ip**) and enter:

ospf-max-paths *<max_number>*

max_number is an integer from 1 (the default) through 5.

For example, the following command sequence enables round-robin ECMP routing and allows up to five distribution paths for RIP and OSPF:

```
ip# ecmp-method roundRobin
ip# rip-max-paths 5
ip# ospf-max-paths 5
```

Using Site Manager

To enable or disable equal-cost multipath support and to specify the maximum number of RIP and OSPF routes, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the following parameters: • **Multiple Nexthop Calculation Method** • **RIP Maximum Equal Cost Paths** • **IP OSPF Maximum Path** Click on **Help** or see the parameter descriptions beginning on page A-25.	
5. Click on **OK**.	You return to the Configuration Manager window.

Customizing the IP Routing Table Structure

Structurally, the IP routing table consists of *indexes* and *entries*. Each index contains a pointer to a list of entries. By default, the IP routing table contains 8000 indexes.

A routing table in which all indexes point to the same number of entries is considered to be in perfect balance. For example, a routing table that contains 100 indexes pointing to 1000 entries is in perfect balance if each index points to 10 entries.

In reality, an IP routing table is allowed to contain indexes that deviate from perfect balance by a number of entries specified as the *deviation-of-nodes* value. By default the deviation-of-nodes value is 25.

Caution: Nortel Networks recommends that you use the default values for the IP routing table parameters. If you want to specify different values, consult the Nortel Networks Technical Solutions Center.

To specify the number of indexes in the IP routing table and to specify a deviation-of-nodes value, navigate to the IP global prompt (for example, **box; ip**) and enter the following commands:

routing-table-indexes *<number>*
routing-table-deviation *<deviation>*

number is the number of indexes in the IP routing table.

deviation is the number of entries by which an index is allowed to deviate from perfect balance.

For example, the following command sequence configures an IP routing table with 1000 indexes and a deviation value of 10:

```
ip# routing-table-indexes 1000
ip# routing-table-deviation 10
```

Specifying the Percentage of Buffers Available to ARP

By default, ARP can use 100 percent of the available buffers for saving buffers when resolving ARP requests. You can reset this percentage.

Using the BCC

To change the percentage of buffers available to ARP, navigate to the IP global prompt (for example, **box; ip**) and enter:

arp-buf-limit *<percentage>*

percentage is the percentage of available buffers that you want to use for resolving ARP requests.

For example, the following command allocates 20 percent of the buffers for resolving ARP requests:

```
ip# arp-buf-limit 20
```

Using Site Manager

To change the percentage of buffers available to ARP, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **Percentage of ARP Buffers** parameter. Click on **Help** or see the parameter description on page A-26.	
5. Click on **OK**.	You return to the Configuration Manager window.

Customizing an IP Interface

An IP network *interface* consists of a physical circuit configured with the appropriate data link and IP protocols. Each interface connects the router to one or more IP networks.

For example, the router in Figure 3-1 is configured with three IP interfaces. One of these interfaces is a point-to-point interface that connects the router to a single long-haul medium terminated by a host or another router. The other two interfaces are LAN interfaces that connect the router to an Ethernet or FDDI local area medium.

An IP interface can provide access to multiple networks. For example, in Figure 3-1, LAN interface 1 provides a connection to both LAN B and LAN C.

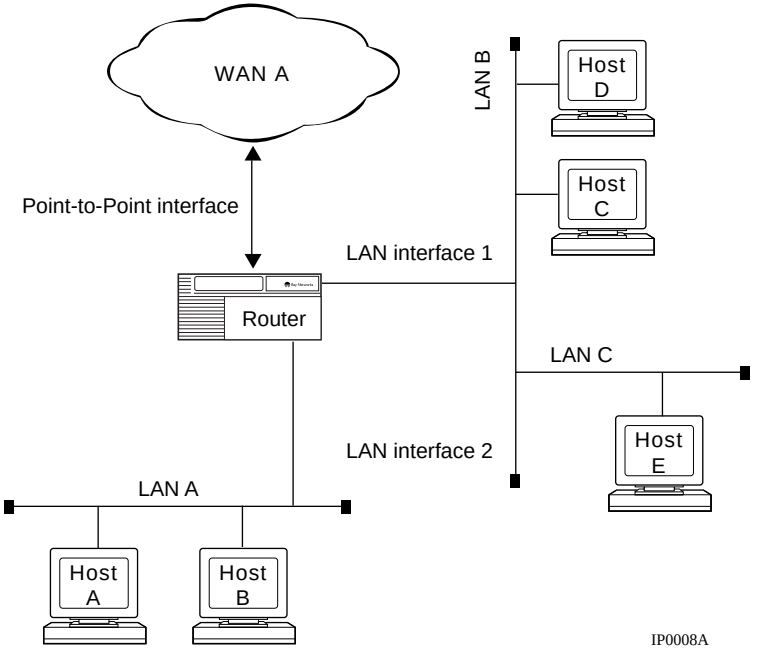

Figure 3-1. IP Interface

When you configure an IP interface on a circuit, the interface is enabled with default values for all interface parameters. You customize an IP interface by modifying parameters as described under the following topics:

Topic	Page
Disabling and Reenabling an IP Interface	3-22
Specifying a Broadcast Address for an Interface	3-23
Specifying the Cost of an Interface	3-24
Enabling MTU Discovery on an Interface	3-26
Enabling and Disabling ICMP Address-Mask Replies	3-27
Disabling and Reenabling ICMP Redirect Messages	3-29
Enabling All-Subnet Broadcasting on an Interface	3-30
Disabling UDP Checksum Processing on an Interface	3-31
Specifying a MAC Address or E.164 Address	3-33
Enabling Source Routing over a Token Ring Network	3-34
Configuring the MTU for an IP Interface	3-37
Configuring an SMDS Address	3-38
Configuring a WAN Address for a Frame Relay Network	3-39
Specifying the Maximum Size of the Forwarding Table	3-40
Configuring an Interface for an ATM Logical IP Subnet	3-42
Deleting IP from an Interface	3-45

Note: When you reconfigure an interface in dynamic mode, IP restarts on that interface. Thus, if the interface you reconfigure is the interface that supports Site Manager's SNMP connection to the router, restarting IP on that interface causes Site Manager to temporarily lose its router connection and to display a warning message. To verify that the change took effect, display the Edit IP Global Parameters window and inspect the settings.

Disabling and Reenabling an IP Interface

When you configure an IP interface on a circuit, the interface is automatically enabled. You can use the BCC or Site Manager to change the state of the IP interface as required.

Using the BCC

To disable or reenable an IP interface, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

state <*state*>

state is one of the following:

enabled (default)
disabled

For example, the following command disables IP interface 2.2.2.2:

```
ip/2.2.2.2# state disabled
ip/2.2.2.2#
```

Using Site Manager

To disable or renable an IP interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-4.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Broadcast Address for an Interface

In broadcasting, the IP router transmits a single packet to every host on an attached network. To do so, it uses a broadcast address that refers to all hosts on the network. A broadcast address is simply an IP address that contains all ones or all zeros in the host portion.

For example, the IP Class C address 10.3.45.12 has the following characteristics:

- Because the address is for a Class A network (the network portion is 1 byte), the host portion contains 3 bytes.
- Because the host portion of a broadcast address consists of all ones or all zeros, the broadcast address for that network can be one of the following: 10.255.255.255, 10.0.0.0, 255.255.255.255, or 0.0.0.0.

You configure a broadcast address for a subnet differently from the way you configure a broadcast address for a network. When you extend the network portion of the IP address to create a subnet address, you automatically take away from the host portion of the address. To configure a subnet broadcast, you take the subnet mask for that subnet and invert it. For example, if the IP address of the subnet is 10.4.2.3, and the mask is 255.255.0.0, then the subnet broadcast address is either 10.4.255.255 or 10.4.0.0.

Some networks do not support broadcasts; thus, configuring an IP broadcast address does not guarantee efficient broadcast delivery.

By default, IP uses a broadcast address that contains all ones in the host portion. Accept the default unless the calculated broadcast address (host portion) of all ones is not adequate. If this is the case, then use the BCC or Site Manager to enter the appropriate IP broadcast address in dotted-decimal notation.

Using the BCC

To specify a broadcast address, navigate to an IP interface prompt (for example, **stack; eth 2/2; ip/2.2.2.2/255.0.0.0**) and enter:

broadcast *<broadcast_address>*

broadcast_address is an IP address expressed in dotted-decimal notation.

For example, the following command assigns broadcast address 1.1.1.1 to IP interface 2.2.2.2.

```
ip/2.2.2.2# broadcast 1.1.1.1
ip/2.2.2.2#
```

Using Site Manager

To specify a broadcast address, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **Broadcast Address** parameter. Click on **Help** or see the parameter description on page A-5.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the Cost of an Interface

Each IP interface has an assigned cost. The interface cost is added to routes learned on this interface through RIP and is specified in subsequent RIP packets transmitted out other interfaces.

If the interface is configured for RIP, keep in mind that increasing the cost causes the upper bound set by the RIP network diameter parameter to be attained more rapidly.

By default, an IP interface has a cost of 1. You can use the BCC or Site Manager to specify another value as required.

Using the BCC

To specify the cost of an IP interface, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

cost <*cost*>

cost is an integer indicating the cost of interface, from 1 through 127.

For example, the following command assigns a cost of 2 to IP interface 2.2.2.2:

```
ip/2.2.2.2# cost 2
ip/2.2.2.2#
```

Using Site Manager

To specify the cost of an IP interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **Cost** parameter. Click on **Help** or see the parameter description on page A-5.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Enabling MTU Discovery on an Interface

A probe MTU is a request for the maximum transmission unit (MTU) size used on all networks that an IP datagram must traverse from source to destination.

By configuring IP to respond to probe MTUs on this interface, you eliminate transit fragmentation and destination reassembly for datagrams destined for this interface and, therefore, decrease network load.

The reply MTU and the probe MTU are options 11 and 12 in RFC 1063.

By default, IP does not respond to probe MTU requests. You can use the BCC or Site Manager to turn this feature on and off as required.

Using the BCC

To enable or disable MTU discovery on an interface, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

mtu-discovery *<state>*

state is one of the following:

on
off (default)

For example, the following command causes IP to respond to probe MTUs on interface 2.2.2.2:

```
ip/2.2.2.2# mtu-discovery on
ip/2.2.2.2#
```

Using Site Manager

To enable or disable MTU discovery on an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **MTU Discovery** parameter. Click on **Help** or see the parameter description on page A-6.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Enabling and Disabling ICMP Address-Mask Replies

You can configure IP to generate Internet Control Message Protocol (ICMP) address-mask reply messages on this interface in response to valid address-mask request messages. The interface generates ICMP address-mask reply messages in compliance with the relevant sections of RFCs 950 and 1009.

By default, IP does not generate address-mask reply messages. You can use the BCC or Site Manager to turn this feature on and off as required.

Using the BCC

To enable or disable ICMP address-mask reply messages, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

mask-reply *<state>*

state is one of the following:

enabled
disabled (default)

For example, the following command causes IP to send ICMP address-mask reply messages on interface 2.2.2.2:

```
ip/2.2.2.2# mask-reply enabled
ip/2.2.2.2#
```

Using Site Manager

To enable or disable ICMP address-mask reply messages, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **AMR** parameter. Click on **Help** or see the parameter description on page A-11.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Disabling and Reenabling ICMP Redirect Messages

An ICMP redirect is a message sent by the router to alert a host that it should be using a different path to route data.

In some cases, you do not want an interface to send out redirects. For example, in a frame relay network, two stations on the same network may not be directly connected if the network is not fully meshed. Thus, in this case, you would disable redirects on this interface.

By default, IP sends ICMP redirect messages. You can use the BCC or Site Manager to disable and reenable this feature on an IP interface as required.

Using the BCC

To disable or reenable ICMP redirect messages, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

redirects *<state>*

state is one of the following:

enabled (default)
disabled

For example, the following command turns off ICMP redirect messages on IP interface 2.2.2.2:

```
ip/2.2.2.2# redirects enabled
ip/2.2.2.2#
```

Using Site Manager

To disable or reenable ICMP redirect messages, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **Redirect** parameter. Click on **Help** or see the parameter description on page A-11.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Enabling All-Subnet Broadcasting on an Interface

An all-subnet broadcast (ASB) datagram has a destination address equal to the broadcast address for an entire network (all subnets). For example, if a network interface serves the subnet 128.10.2.1 with a subnet mask of 255.255.255.0, the IP router considers any datagram with a destination address of 128.10.255.255 or 128.10.0.0 to be an ASB datagram.

By default, IP does not flood ASB datagrams. You can use the BCC or Site Manager to turn this feature on and off as required.

Using the BCC

To enable or disable all-subnet broadcasting, navigate to an IP interface prompt (for example, **stack; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

all-subnet-broadcast *<state>*

state is one of the following:

enabled
disabled (default)

For example, the following command causes IP to flood ASB datagrams out interface 2.2.2.2:

```
ip/2.2.2.2# all-subnet-broadcast enabled
ip/2.2.2.2#
```

Using Site Manager

To enable or disable all-subnet broadcasting, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **ASB** parameter. Click on **Help** or see the parameter description on page A-6.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Disabling UDP Checksum Processing on an Interface

By default, UDP checksum processing is enabled on this interface. All outgoing and incoming UDP datagrams are subject to checksum processing. You can use the BCC or Site Manager to turn this feature on and off as required. Use the default in virtually all instances. Disable UDP checksum processing to provide backward compatibility with UNIX BSD 4.1.

Using the BCC

To disable or reenable UDP checksum processing, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

udp-checksum <*state*>

state is one of the following:

enabled (default)
disabled

For example, the following command turns off UDP checksum processing on IP interface 2.2.2.2:

```
ip/2.2.2.2# udp-checksum disabled
ip/2.2.2.2#
```

Using Site Manager

To disable or reenable UDP checksum processing, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **Udp Xsum On** parameter. Click on **Help** or see the parameter description on page A-9.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a MAC Address or E.164 Address

You can use Site Manager or the BCC to specify a MAC address or an E.164 address for this interface. The IP router uses its address and this MAC address when transmitting and receiving packets on this interface. Table 3-3 lists the valid settings for the MAC address parameter.

Table 3-3. MAC Address Parameter Settings

Value	Meaning
0	The IP router uses its IP address and the circuit's MAC address when transmitting packets on this interface.
User-specified MAC address	The IP router uses its IP address and this MAC address when transmitting and receiving packets on this interface.
E.164 address	If the interface is on an SMDS circuit, by default IP uses the individual SMDS-configured address. You can enter the entire E.164 address -- for example, C1 617 555 5000 FFFF.

Using the BCC

To specify a MAC address, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

configured-mac-address *<address>*

address is one of the following:

0
User-specified MAC address

No default value exists for this parameter.

For example, the following command directs the router to use its own IP address and the circuit's MAC address on IP interface 2.2.2.2:

```
ip/2.2.2.2# configured-mac-address 0
ip/2.2.2.2#
```

Using Site Manager

To configure this parameter for a multinet or multigroup configuration, see *Configuring SMDS.*

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **MAC Address** parameter. Click on **Help** or see the parameter description on page A-9.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Enabling Source Routing over a Token Ring Network

The IP router can route over token ring networks that contain one or more source routing bridges.

In a source routing network, every end station that sends a frame supplies the frame with the necessary route descriptors so that it can be source routed across the network. Thus, in order for an IP router to route packets across a source routing network, the router must act as an end station, supplying route descriptors for each packet before it sends it onto the network.

With end-station support enabled, an IP router does the following (Figure 3-2):

1. Receives a packet and determines that the packet's next hop is located across a source routing network
2. Adds the necessary routing information field (RIF) information to the packet's MAC header
3. Sends the packet onto the network where it is source routed toward the next hop

Upon receiving the packet from the token ring network, the peer router strips off the RIF and continues to route the packet toward the destination network address.

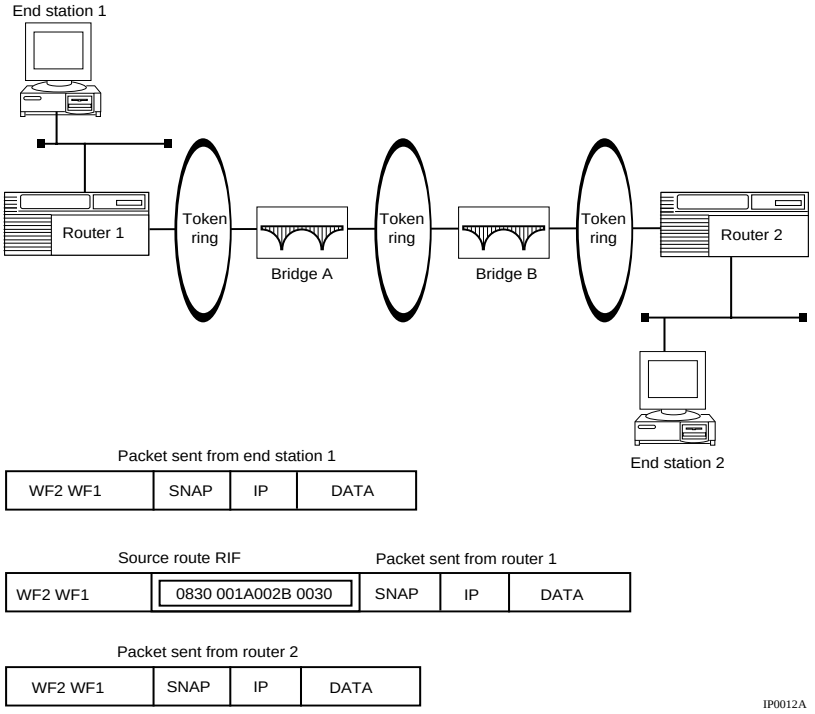

Figure 3-2. IP Routers Source Routing Across a Token Ring Network

The router can send ARP packets over an interface configured for a token ring network. Nortel Networks supports both spanning tree explorer (STE) and all routes explorer (ARE) ARP packets.

Using the BCC

To disable or reenable source route end-station support, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

end-station-support *<state>*

state is one of the following:

enabled
disabled (default)

For example, the following command activates source route end-station support on IP interface 2.2.2.2:

```
ip/2.2.2.2# end-station-support enabled
ip/2.2.2.2#
```

Using Site Manager

To configure source route end-station support on a per-circuit basis and choose STE or ARE ARP packets, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the following parameters: • **TR End Station** • **TR Endstation ARP Type** Click on **Help** or see the parameter descriptions beginning on page A-10.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring the MTU for an IP Interface

You can set the MTU for an interface to a value less than the MTU of the underlying circuit. The default value (0) causes the IP interface to use the MTU value of the underlying circuit. If you configure a value greater than the MTU of the underlying circuit, the parameter is ignored.

Setting the MTU for an interface affects the transmission of IP frames only. Transmitted frames that are larger than the interface MTU value are fragmented into smaller frames. This parameter does not affect frames accepted by the driver.

Using the BCC

To set the MTU for the IP interface, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

mtu *<max_units>*

max_units is the MTU for this IP interface.

For example, the following command sets the MTU to 2000 on interface 2.2.2.2:

```
ip/2.2.2.2# mtu 2000
ip/2.2.2.2#
```

Using Site Manager

To set the MTU for the IP interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **Interface MTU** parameter. Click on **Help** or see the parameter description on page A-10.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring an SMDS Address

By default, if the interface is connected to an SMDS network, IP uses the SMDS-configured addresses.

You can use Site Manager to supply the following:

- A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider
- An address-resolution multicast address for this IP interface in an SMDS network

Note: Make sure to enter the correct addresses for the MAC Address, SMDS Group Address, and SMDS Arp Request Address parameters. These addresses are the same as those you entered for the Individual Address, Group Address, and ARP Address parameters in the SMDS Configuration window when you configured SMDS.

To configure an SMDS address, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the following parameters: • **SMDS Group Address** • **SMDS Arp Request Address** Click on **Help** or see the parameter descriptions beginning on page A-12.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a WAN Address for a Frame Relay Network

If an interface is connected to a frame relay network, you can use Site Manager to configure the following:

- A broadcast address. If you enter a value for the FRM Broadcast parameter, the frame relay switch, rather than the router, will broadcast the message.
- A multicast address for this IP interface that will send messages to all OSPF routers in a frame relay network. If you enter a value for the FRM Cast 1 DLCI parameter, the frame relay switch, rather than the router, will send the message to all OSPF routers.
- A multicast address for this IP interface that will send messages to all OSPF designated routers in a frame relay network. If you enter a value for the FRM Cast 2 DLCI parameter, the frame relay switch, rather than the router, will send the message to all OSPF designated routers.

To configure broadcast and multicast addresses to be used by the frame relay switch, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the following parameters: • **FRM Broadcast** • **FRM Cast 1 DLCI** • **FRM Cast 2 DLCI** Click on **Help** or see the parameter descriptions beginning on page A-13.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the Maximum Size of the Forwarding Table

To minimize the amount of time it spends looking up routes, IP creates and maintains a cache of frequently used routes -- called a *forwarding table* or *cache* -- for each IP interface. A forwarding table is a first in first out (FIFO) buffer.

When a datagram arrives on an IP interface for forwarding, IP searches the forwarding table associated with the interface for the destination network. If the search is successful, IP dispatches the datagram to the interface noted in the table entry.

If the search is unsuccessful, IP consults the routing table to obtain the same information, dispatches the datagram to the appropriate interface, and caches the information in the appropriate forwarding table -- either by appending information to the table (if the table is not full) or by overwriting the oldest, first-in table entry (if the table is full).

When IP flushes a route from the routing table, it also removes the route from the forwarding tables, thus ensuring that invalid routing information is not retained in interface-specific caches.

An interface that receives packets that are destined for a large number of different destinations may benefit from a larger forwarding table. The larger the number of entries, the more likely it is that the destination will already be in the forwarding table and the faster the route lookups will be for those destinations.

Keep in mind that configuring a forwarding table size that is larger than necessary reduces the total amount of memory usable by other applications. On the other hand, configuring a forwarding table size that is too small can affect overall router performance.

Check the number of cache hits and misses to determine the optimal size of the forwarding table. For debugging purposes, if you see the wfIpInterfaceCacheMisses statistic going up at a rapid rate, consider increasing the table size. However, an occasional cache miss does not warrant an increase in the table size.

By default, IP allocates a cache for 128 destination entries on the interface. You can specify a different cache size.

Using the BCC

To set the maximum number of entries allowed in the forwarding table, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

cache-size *<size>*

size is the number of destination entries in the cache.

For example, the following command causes IP to allocate a cache on interface 2.2.2.2 for 175 entries:

```
ip/2.2.2.2# cache-size 175
ip/2.2.2.2#
```

Using Site Manager

To set the maximum number of entries allowed in the forwarding table, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **Forward Cache Size** parameter. Click on **Help** or see the parameter description on page A-16.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring an Interface for an ATM Logical IP Subnet

RFC 1577, “Classical IP and ARP over ATM,” is a specification for an administrative entity within an ATM network called a logical IP subnet (LIS). Nortel Networks supports RFC 1577. For information about configuring IP interfaces on an ATM LIS, see *Configuring ATM Services.*

You can use the BCC or Site Manager to do the following:

- Specify the ATMARP mode: client or server. You must configure one ATMARP server for each logical IP subnet that you define.
- Define the ATM address network prefix of the ATMARP server on your network. A complete ATM address consists of a network prefix and a user part.
- Define the user part (suffix) of the ATM address for the ATMARP server on your network. The user part consists of a 6-byte end-station identifier and a 1-byte selector field.
- Specify (for a client) the interval between registration refreshes.

- Specify (for a server) the duration for which the registration is valid.

Using the BCC

To set parameters for an ATM logical IP subnet, navigate to an IP interface prompt on a classical IP service record (for example, **box; atm/11/1; classical-ip-service/paris; ip/2.2.2.2/255.255.0.0**) and enter:

arp-mode *<mode>*

mode is one of the following:

client (default)
server

To specify the ATM address of an ATMARP server, enter:

atmarp-server-address *<address>*

address is an ATM hexadecimal address (40 characters).

To specify a server registration interval, enter:

atmarp-server-interval *<interval>*

interval is any integer. The default value for a client is 900 seconds; for a server, 1200 seconds.

To disable or reenable timing out of the connection to the ATMARP server, enter:

atmarp-server-aging *<state>*

state is either:

enabled (default)
disabled

Using Site Manager

To set parameters for an ATM logical IP subnet, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the following parameters: • **ATM ARP Mode** • **ARP Server ATM Address Network Prefix** • **ARP Server ATM Address User Part** • **Registration Refresh Interval** Click on **Help** or see the parameter descriptions beginning on page A-17.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Deleting IP from an Interface

You can delete IP from an interface using the BCC or Site Manager.

Using the BCC

To delete IP from an interface, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

delete

For example, the following command deletes IP from interface 2.2.2.2:

```
ip/2.2.2.2/255.255.0.0# delete
ethernet/2/1#
```

Using Site Manager

To delete IP from an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on the connector from which you want to delete IP services.	The Edit Connector window opens.
2. Click on **Edit Circuit**.	The Circuit Definition window opens.
3. Choose **Protocols**.	The Protocols menu opens.
4. Choose **Add/Delete**.	The Select Protocols window opens. The IP button is checked to show that IP is enabled on the circuit.
5. Click on **IP**.	
6. Click on **OK**.	Site Manager deletes IP services from the connector and you return to the Circuit Definition window.
7. Choose **File**.	The File menu opens.
8. Choose **Exit**.	You return to the Configuration Manager window.

Configuring an Adjacent Host Address

An *adjacent host* is a device on a locally attached network. This device may or may not be a router. You must configure a MAC address for each adjacent host that does not implement ARP.

Also, if a local network does implement ARP, you may want to configure a MAC address for an adjacent host to preempt the ARP process.

You can use the BCC or Site Manager to configure an adjacent host.

Using the BCC

To associate the IP address of an adjacent host with its physical address, navigate to the global IP prompt (for example, **box; ip**) and enter one of the following:

adjacent-host ip-address *<ip_address>* **mac-address** *<physical_address>*

adjacent-host ip-address *<ip_address>* **vpi-vci** *<physical_address>*

adjacent-host ip-address *<ip_address>* **nsap** *<physical_address>*

adjacent-host ip-address *<ip_address>* **dlci** *<physical_address>*

adjacent-host ip-address *<ip_address>* **wan-address** *<physical_address>*

ip_address is the address of the adjacent host.

physical_address is the address that you want to associate with the IP address (see Table 3-4).

The BCC displays a prompt for the adjacent host. To set parameters that define the adjacent host, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 3-4.

Table 3-4. BCC Adjacent Host Parameters

Parameter	Value/Default	Description
state	enabled (default) disabled	Specifies the state of the adjacent host definition
ip-address	No default	Specifies the IP address of the device for which you want to configure an adjacent host
mac-address	No default	Specifies the physical address of the adjacent host. Enter the MAC address as a 12-digit hexadecimal number.
vpi-vci	No default	Specifies the physical address of the adjacent host. Enter an ATM PVC address in the form virtual path identifier/virtual channel identifier, for example, 0/32.
nsap	No default	Specifies the physical address of the adjacent host
dlci	No default	Specifies the physical address of the adjacent host
wan-address	No default	Specifies the physical address of the adjacent host
encapsulation	ethernet (default) snap null	Specifies the adjacent host's encapsulation method. Select ethernet or snap (Service Network Access Point) if you are defining a point-to-point network interface or if the adjacent host resides on an Ethernet. For an adjacent host on an ATM logical IP subnet, select snap or null.
type	default (default) e164 x121	Specifies the type of adjacent host
sub-address	No default	Specifies the subaddress used to establish an SVC to the adjacent host
type-of-number	international (default) unknown	Specifies the type of number used to establish an SVC to the adjacent host

Using Site Manager

To associate the IP address of an adjacent host with its physical address, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Adjacent Hosts**.	The IP Adjacent Hosts window opens.
4. Click on **Add**.	The IP Configuration window opens.
5. Set the following parameters: • **IP Adjacent Host Address** • **MAC/ATM Address, DLCI, VPI/VCI** • **Host Encapsulation** • **Adjacent Host Address** • **Remote Party Sub-Address** • **Remote Party Type of Number** • **Adjacent Host Type** Click on **Help** or see the parameter descriptions beginning on page A-30.	
6. Click on **OK**.	The IP Adjacent Hosts window displays the adjacent host you just configured.
7. Set the following parameters: • **Enable** • **Next Hop Interface Addr** Click on **Help** or see the parameter descriptions beginning on page A-30.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Defining a Static Route

A *static route* is a manually configured route that specifies the transmission path a datagram must follow, based on the datagram's destination address. A static route specifies a transmission path to another network. You configure a static route if you want to restrict datagrams to paths you specifically configure.

Static routes remain in IP routing tables until you remove them. Note, however, that if the interface that was used to reach the next hop in the static route becomes disabled, the static route is deleted from the IP routing table.

You can use the BCC or Site Manager to specify the following:

- The state (active or inactive) of the static route record in the IP routing tables.
- The IP address of the network to which you want to configure the static route.
- The subnet mask of the destination network.
- The number of router hops that a datagram can traverse before reaching the destination IP address. The IP router uses the cost value when determining the best route for a datagram to follow.
- The IP address of the next-hop router.
- The subnet mask of the next-hop router.
- A weighted value (from 1 to 16, with 16 being the most preferred) that the IP router uses to choose a route when its routing tables contain multiple routes to the same destination.
- The local router circuit associated with the static route over an unnumbered interface.

IP supports multiple static routes to the same destination. IP uses the best route to forward packets and treats the other routes as backup routes in case the chosen route becomes unusable or is no longer considered the best route.

You can also configure IP to support equal-cost multipath (ECMP) routes for traffic load balancing. If IP considers the ECMP routes to be the best routes, IP uses them all in the way you specify -- in round-robin fashion, for example -- to forward data. For information, see "Configuring Equal-Cost Multipath Support for RIP and OSPF" on page 3-15. With ECMP enabled globally on the router, you can configure up to 12 ECMP static routes.

Note: When you dynamically create a new static route (for example, 128.10.8.0/255.255.255.0/1), the router may generate the following error message in the router log file:

```
# 2: 04/16/99 09:44:24.096 WARNING SLOT 5 IP Code: 56
Error when configuring static route 128.10.8.0/255.255.255.0/1
```

This message indicates that the static route is not configured at the time stated because all necessary MIB sets are not completed. Error checking occurs as each attribute is set. The router code generates the error messages to prevent misconfigured static routes. You can verify that a static route has been created correctly by issuing a **get** command for the valid attribute in the wfIpStaticRouteEntry.10.* instance.

Using the BCC

To define a static route, navigate to the global IP prompt (for example, **box; ip**) and enter:

static-route address *<destination>* **mask** *<ip_mask>* **next-hop-address** *<next_hop>*

destination is the destination IP address.

ip_mask is the mask of the destination IP address.

next_hop is the next-hop IP address.

The static route prompt appears.

The BCC configures a static route with default values for all static route parameters and displays a static-route-specific prompt. You customize a static route by modifying static route parameters. Navigate to the static route prompt and enter:

<parameter> <value>

parameter value is one of the parameter/value pairs described in Table 3-5.

For example, the following command line configures a static route to destination 3.2.4.5/255.255.0.0:

```
ip# static-route address 3.2.4.5 mask 255.255.0.0 next-hop-address 2.2.2.4
static-route/3.2.4.5/255.255.0.0/2.2.2.4#
```

Table 3-5. BCC Static Route Parameters

Parameter	Values/Defaults	Meaning/Instructions
state	enabled (default) disabled	Specifies the state (active or inactive) of the static route record in the IP routing tables. Select disabled to make the static route record inactive in the IP routing table; the IP router will not consider this static route. Select enabled to make the static route record active again in the IP routing table.
address	The destination IP address you supplied when you configured the static route. You cannot modify this address.	Specifies the IP address of the network to which you want to configure the static route. Enter the destination IP address in dotted-decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet address. You can configure up to 12 static routes to the same destination.
mask	The destination IP address you supplied when you configured the static route. You cannot modify this address.	Specifies the subnet mask of the destination network. Enter the subnet or supernet mask in dotted-decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet mask.
next-hop-address	The next-hop IP address you supplied when you configured the static route. You cannot modify this address.	Specifies the IP address of the next-hop router. Enter the IP address in dotted-decimal notation. To configure a black hole static route, enter 255.255.255.255. If you are configuring a static route to an unnumbered interface, enter 0.0.0.0.
next-hop-mask	IP address mask	Specifies the subnet mask of the next-hop router.

(continued)

Table 3-5. BCC Static Route Parameters *(continued)*

Parameter	Values/Defaults	Meaning/Instructions
cost	1 (default) to RIP diameter	Specifies the number of router hops a datagram can traverse before reaching the destination IP address. The IP router uses the cost value when determining the best route for a datagram to follow. If you enabled ECMP on the router (for information, see "Configuring Equal-Cost Multipath Support for RIP and OSPF" on page 3-15), you can configure up to 12 equal-cost static routes.
preference	1 (default) to 16	Specifies a weighted value (from 1 to 16, with 16 being the most preferred) that the IP router uses to select a route when its routing tables contain multiple routes to the same destination. To configure a black hole static route, enter the maximum preference value.

Using Site Manager

To define a static route, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Static Routes**.	The IP Static Routes window opens.
4. Click on **Add**.	The IP Configuration window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the following parameters: • **Destination IP Address** • **Address Mask** • **Cost** • **Next Hop Addr** • **Next Hop Mask** • **Preference** • **Unnumbered CCT Name** Click on **Help** or see the parameter descriptions beginning on page A-27.	
6. Click on **OK**.	You return to the IP Static Routes window.

Defining a Static Default Route

If IP receives a data packet with a destination address that it is unable to match in its routing table, it looks for a default route that it can use to forward the packet.

To include a default route in the routing table, create a static route with a destination address of 0.0.0.0. For the next-hop address, specify a router that can forward the packet to its destination.

Defining a Static Black Hole for a Supernet

A router that advertises an aggregate route by using a supernet address to represent multiple explicit routes must be able to discard packets that match the supernet address but that do not match any of the explicit routes.

For example, consider a router that advertises an aggregate route using the supernet address 192.32.0.0/255.255.255.248. The supernet address represents eight specific networks: 192.32.0.0 through 192.32.7.0. After the aggregate route has been propagated, the router receives network traffic for each of these specific destinations.

If the router loses connectivity to network 192.32.3.0, one of the networks in the supernet, the router continues to forward traffic that matches destinations 0.0 through 2.0 and 4.0 through 7.0. However, the router can no longer find a complete match in the routing table for the disconnected network, 3.0. The router must drop all traffic destined for 192.32.3.0.

To force the router to drop the packet for an unmatched destination, you configure a special type of static route for a supernet called a *black hole*, as follows:

- Enter the supernet address/mask pair as the destination IP address and address mask.
- To create the black hole, enter the black hole encoding (255.255.255.255) as the next-hop address and the next-hop mask.

Configuring and Customizing Router Discovery

Before a host can send IP datagrams beyond its directly attached subnet, the host must discover the address of at least one operational router on that subnet. Router Discovery is an extension of the Internet Control Message Protocol (ICMP) that enables hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers.

Routers configured with Router Discovery periodically multicast or broadcast a router advertisement from each of their interfaces, announcing the IP address or addresses of that interface. Hosts discover the addresses of their neighboring routers by listening for these advertisements. Hosts will use the router with the highest preference level as a gateway.

By default, Router Discovery is enabled on each IP interface. You can use Site Manager to choose the operating characteristics of Router Discovery on the interface, as described under the following topics:

Topic	Page
Enabling and Disabling Router Discovery	3-55
Choosing a Broadcast Type	3-56
Specifying the Time Interval Between Advertisements	3-57
Configuring the Lifetime of Advertised Addresses	3-58
Specifying Interface Preference	3-59

Enabling and Disabling Router Discovery

You can use the BCC or Site Manager to enable and disable Router Discovery on an interface.

Using the BCC

To enable or disable Router Discovery on an interface, navigate to a Router Discovery prompt on an IP interface (for example, **box; eth 2/2; ip/2.2.2.2/ 255.255.0.0; rdisc**) and enter:

state <*state*>

state is one of the following:

enabled (default)
disabled

For example, the following command disables Router Discovery on interface 2.2.2.2:

```
rdisc/2.2.2.2# state disabled
rdisc/2.2.2.2#
```

Using Site Manager

To enable or disable Router Discovery on an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Router Discovery**.	The IP Router Discovery window opens.
4. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-33.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Choosing a Broadcast Type

You can use the BCC or Site Manager to specify the type of broadcast to use in sending advertisements. You should use multicast (the default broadcast type) wherever possible; that is, on any link where all listening hosts support IP multicast.

Using the BCC

To specify a broadcast type for sending advertisements, navigate to a Router Discovery prompt on an IP interface (for example, **box; eth 2/2; ip/2.2.2.2/ 255.255.0.0; rdisc**) and enter:

mode *<state>*

state is one of the following:

mcast (default)
local
direct

For example, the following command sets the broadcast type to local for Router Discovery on interface 2.2.2.2:

```
rdisc/2.2.2.2# mode local
rdisc/2.2.2.2#
```

Using Site Manager

To specify a broadcast type for sending advertisements, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Router Discovery**.	The IP Router Discovery window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Set the **Broadcast Type** parameter. Click on **Help** or see the parameter description on page A-33.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the Time Interval Between Advertisements

You can use the BCC or Site Manager to specify the minimum and maximum number of seconds between advertisements. The default values for the minimum and maximum intervals are 450 and 600 seconds respectively.

Using the BCC

To specify the minimum and maximum number of seconds between advertisements, navigate to a Router Discovery prompt on an IP interface (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rdisc**) and enter:

min-interval *<interval>* **max-interval** *<interval>*

interval is any integer.

For example, the following command sets the minimum and maximum intervals to 550 and 750 respectively on interface 2.2.2.2:

```
rdisc/2.2.2.2# min-interval 550 max-interval 750
rdisc/2.2.2.2#
```

Using Site Manager

To specify the minimum and maximum number of seconds between advertisements, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Router Discovery**.	The IP Router Discovery window opens.
4. Set the following parameters: • **Minimum Interval** • **Maximum Interval** Click on **Help** or see the parameter descriptions beginning on page A-33.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring the Lifetime of Advertised Addresses

You can use the BCC or Site Manager to specify the maximum length of time that advertised addresses are to be considered as valid router addresses by hosts, in the absence of further advertisements. The default value is 1800 seconds.

Using the BCC

To specify the maximum length of time that advertised addresses are considered valid by hosts, navigate to a Router Discovery prompt on an IP interface (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rdisc**) and enter:

lifetime *<interval>*

interval is any integer.

For example, the following command sets the lifetime of advertised addresses to 2500 seconds on interface 2.2.2.2:

```
rdisc/2.2.2.2# lifetime 2500
rdisc/2.2.2.2#
```

Using Site Manager

To specify the maximum length of time that advertised addresses are considered valid by hosts, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Router Discovery**.	The IP Router Discovery window opens.
4. Set the **Lifetime** parameter. Click on **Help** or see the parameter description on page A-34.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying Interface Preference

You can use the BCC or Site Manager to specify the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The default value is 0.

Using the BCC

To specify the preference value of the address as a default router address, navigate to a Router Discovery prompt on an IP interface (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rdisc**) and enter:

preference *<integer>*

integer is the metric value.

For example, the following command sets the preference value to 1 on interface 2.2.2.2:

```
rdisc/2.2.2.2# preference 1
rdisc/2.2.2.2#
```

Using Site Manager

To specify the preference value of the address as a default router address, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Router Discovery**.	The IP Router Discovery window opens.
4. Set the **Interface Preference** parameter. Click on **Help** or see the parameter description on page A-34.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Chapter 4
Customizing ARP and RARP Services

You can customize the Address Resolution Protocol (ARP) and Reverse Address Resolution Protocol (RARP) by setting parameters as described in the following sections:

Topic	Page
Customizing ARP	4-2
Customizing RARP	4-12

Customizing ARP

You can customize the Address Resolution Protocol (ARP) by setting parameters as described under the following topics:

Topic	Page
Enabling and Disabling Global ARP	4-2
Customizing Global ARP	4-3
Selecting an Address Resolution Scheme for an IP Interface	4-5
Selecting an Encapsulation Option for ARP and Probe	4-7
Enabling Proxy ARP on an Interface	4-8
Timing Out Entries in the Address Resolution Cache	4-10

Enabling and Disabling Global ARP

ARP is configured and enabled on the router at startup. You can use the BCC to disable and reenable ARP as required.

To disable or reenable ARP, navigate to the global IP prompt (for example, **box; ip**) and enter:

arp

The global ARP prompt appears.

Enter:

state *<state>*

state is one of the following:

enabled (default)
disabled

For example, the following command sequence disables ARP on the router:

```
ip# arp
arp# state disabled
arp#
```

Customizing Global ARP

You can use the BCC or Site Manager to do the following:

- Control how ARP acts in relation to IP's forwarding state.
- Control whether IP drops and logs an invalid ARP source address or simply drops the request.
- Control whether IP drops or accepts ARP requests in which the source and destination addresses are located in different networks or subnetworks. This parameter allows Proxy ARP to generate replies when the source and destination networks in the ARP request differ.

Using the BCC

To customize global ARP, navigate to the ARP global prompt (for example, **stack; ip; arp**) and enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 4-1.

Table 4-1. BCC Global ARP Parameters

Parameter	Values	Function
nonlocal-source	drop (default) accept	Specifies whether IP drops or accepts ARP requests in which the source address is located in a different network or subnetwork.
nonlocal-destination	drop (default) accept	Specifies whether IP drops or accepts ARP requests in which the destination address is located in a different network or subnetwork.

For example, the following command enables IP to accept ARP requests from nonlocal source addresses:

```
arp# nonlocal-source accept
arp#
```

Using Site Manager

To customize global ARP, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the following parameters: • **ARP Forwarding** • **Nonlocal ARP Source** • **Nonlocal ARP Destination** Click on **Help** or see the parameter descriptions beginning on page A-20.	
5. Click on **OK**.	You return to the Configuration Manager window.

Selecting an Address Resolution Scheme for an IP Interface

In addition to ARP, IP supports Inverse ARP, HP Probe, and X.25 address-resolution schemes as follows:

- Inverse ARP provides address resolution for frame relay interfaces. Use Inverse ARP to discover the IP address of the station at the remote end of the virtual circuit.
- HP Probe, a Hewlett-Packard proprietary protocol, is an address resolution mechanism that functions much like ARP to determine a network host's physical address using the host's IP address by binding a 32-bit IP address to a 48-bit MAC address. IP supports HP Probe over Ethernet and the following HP Probe messages:
  - -- Unsolicited Reply (incoming and outgoing)
  - -- Name Request (incoming)
  - -- Name Reply (outgoing)
  - -- Virtual Address Reply (incoming and outgoing)
  - -- Virtual Address Request (incoming and outgoing)
  - -- Proxy Request (incoming and outgoing)
  - -- Proxy Reply (incoming and outgoing)

  IP can support the concurrent operation of HP Probe and ARP on an interface.

Note: If bridging is configured and enabled on the interface (in addition to IP), the Name Request/Reply and the Proxy Request/Reply messages are bridged.

- ATMARP provides address resolution for a logical IP subnet (LIS) within an ATM network, as defined in RFC 1577, *Classical IP and ARP over ATM*. A router that needs the ATM address of a host on the LIS sends an ATMARP request to the ATMARP server. For more information about classical IP over ATM, see *Configuring ATM Services*.
- The X.25 address-resolution scheme is used on network interfaces that support the X.25 DDN service.
- The RFC 877-compliant address-resolution mechanism is used on network interfaces that support the X.25 PDN service.

On interfaces configured for a token ring network, the router can send ARP requests as spanning tree explorer (STE) or all-routes explorer (ARE) packets.

By default, ARP is enabled on the interface. You can use the BCC or Site Manager to specify an address-resolution scheme.

Using the BCC

To specify the address-resolution mechanism, navigate to an IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

address-resolution *<type>*

type is one of the following:

arp (default)
ddn
pdn
inarp
arpinarp
none
bfeddn
probe
arpprobe
atmarp

Using Site Manager

To specify the address-resolution mechanism, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **Address Resolution Type** parameter. Click on **Help** or see the parameter description on page A-7.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Selecting an Encapsulation Option for ARP and Probe

If you select ARP, Probe, or ARP/Probe, you must also select the appropriate data link encapsulation option as follows:

- If your address-resolution scheme is ARP only, select Ethernet encapsulation, SNAP encapsulation, or Ethernet/SNAP encapsulation.
- If your resolution scheme is HP Probe only, select LSAP encapsulation.
- If your resolution scheme is ARP/Probe, select Ethernet/LSAP encapsulation, SNAP/LSAP encapsulation, or Ethernet/SNAP/LSAP encapsulation.

IP ignores this parameter if the underlying media is anything other than Ethernet.

By default, IP uses ARP Ethernet encapsulation. You can use Site Manager to specify an encapsulation scheme.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **Ethernet Arp Encaps** parameter. Click on **Help** or see the parameter description on page A-12.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Enabling Proxy ARP on an Interface

Proxy ARP allows a router to answer a local ARP request for a remote destination. For example, in Figure 4-1, hosts B and C are located on the same network but on separate subnetworks. Hosts B and C do not understand subnetworking. The router connecting the two physical networks knows which host resides on which network. The address mask is 255.255.255.000. In this example, one subnet is a remote network with respect to the other subnet.

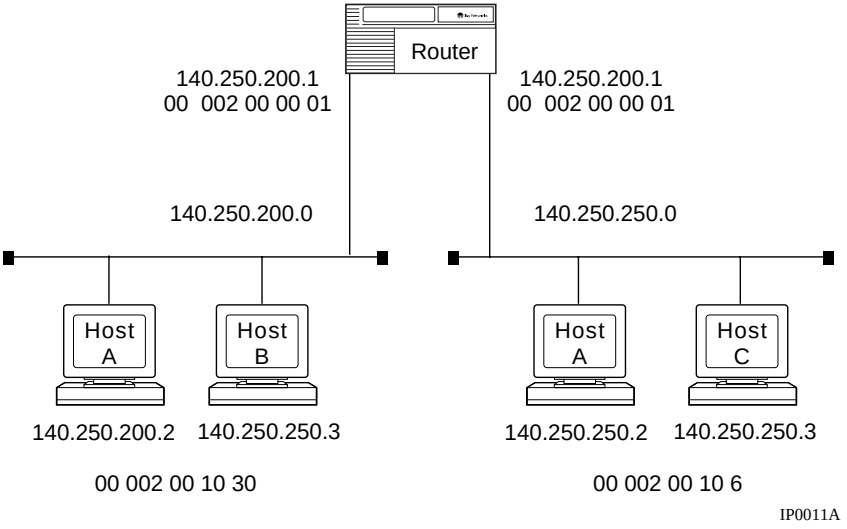

Figure 4-1. Proxy ARP Example

Host B wants to talk to host C, so host B broadcasts an ARP request, which asks IP address 140.250.250.3 to respond with its physical address. The router captures host B's ARP request and responds with its hardware address, 00 002 00 00 01, and host C's IP address, 140.250.250.3. Host B maps host C's IP address (140.250.250.3) to the router's hardware address (00 002 00 00 01).

With Proxy ARP enabled, the router responds with an ARP reply if there is a valid route (that is, if the router is able to forward traffic) to the destination in the routing table. This route may be a subnet route or a default route. For the router to respond for subnets that are reachable via the default route, you must configure IP to use a default route for unknown subnets (see "Using a Default Route for an Unknown Subnet" on page 3-12).

Some devices use Proxy ARP to determine a gateway rather than relying on a statically defined default gateway. These devices use ARP for all remote destinations. To enable the router to reply to ARP for remote destinations on other networks, you must enable Proxy ARP and set the Nonlocal ARP Destination parameter to Accept (see "Customizing Global ARP" on page 4-3).

By default, Proxy ARP is disabled on the interface. You can use the BCC or Site Manager to enable Proxy ARP.

Using the BCC

To enable or disable Proxy ARP, navigate to the IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

proxy *<state>*

state is one of the following:

enabled
disabled (default)

For example, the following command enables Proxy ARP on IP interface 2.2.2.2:

```
ip/2.2.2.2# proxy enabled
ip/2.2.2.2#
```

Using Site Manager

To enable or disable Proxy ARP, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.
5. Set the **Proxy** parameter. Click on **Help** or see the parameter description on page A-8.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Timing Out Entries in the Address Resolution Cache

IP maintains an address-resolution cache on each interface that is configured with ARP or Proxy ARP. The address-resolution cache contains host physical addresses learned by means of ARP or Proxy ARP.

If you enable the cache timeout feature on this interface, the IP router removes address-resolution cache entries after the specified aging period has elapsed. After an entry is removed, the IP router must use ARP to reacquire the physical layer address.

By default, the cache timeout feature is disabled on the interface. You can use the BCC or Site Manager to enable the feature and to specify a timeout interval in seconds.

Using the BCC

To enable or disable the ARP cache timeout feature, navigate to the IP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0**) and enter:

host-cache-aging *<action>*

action is one of the following:

cache-off (default)
cache120
cache180
cache240
cache300
cache600
cache900
cache1200

For example, the following command causes IP to time out entries from the address-resolution cache on IP interface 2.2.2.2 after 300 seconds:

```
ip/2.2.2.2# host-cache-aging cache300
ip/2.2.2.2#
```

Using Site Manager

To enable or disable the ARP cache timeout feature, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Interfaces**.	The IP Interface List window opens.
4. Click on the interface that you want to edit.	Site Manager displays the parameter values for that interface in the IP Interface List window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **Host Cache** parameter. Click on **Help** or see the parameter description on page A-8.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Customizing RARP

You can customize the Reverse Address Resolution Protocol (RARP) by setting parameters as described under the following topics:

Topic	Page
Disabling and Enabling RARP on an Interface	4-13
Defining the RARP Mapping Table	4-13
Disabling and Reenabling RARP Globally	4-14
Deleting RARP Globally	4-15

Disabling and Enabling RARP on an Interface

You can disable or reenable individual RARP interfaces by selecting them from the list of interfaces in the RARP Interface Table.

To disable or reenable RARP on an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Reverse ARP**.	The Reverse ARP menu opens.
4. Choose **Interface Table**.	The RARP Interface Table window opens.
5. Set the **Enable/Disable** parameter. Click on **Help** or see the parameter description on page A-35.	
6. Click on **Apply**.	
7. Click on **Done**.	You return to the Configuration Manager window.

Defining the RARP Mapping Table

The RARP mapping table lists the clients on the network that use the router's RARP services. You define the router's MAC address-to-IP address RARP mapping table as follows:

- Specify the MAC address of each client that will use the RARP services of this router. The client will include the MAC address that you specify in RARP broadcasts to the router.
- Specify the IP address corresponding to the value of the MAC address that you specify. Do not accept the default value (0.0.0.0). When the router receives a RARP request from the client, it assigns this IP address to the client and includes it in a response packet.

To specify a client's MAC address and its IP address, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Reverse ARP**.	The Reverse ARP menu opens.
4. Choose **Map Table**.	The RARP Map Table window opens.
5. Click on **Add**.	The RARP Addresses window opens.
6. Set the following parameters: • **MAC Address** • **IP Address** Click on **Help** or see the parameter descriptions beginning on page A-35.	
7. Click on **OK**.	You return to the RARP Map Table window.
8. Click on **Done**.	You return to the Configuration Manager window.

Disabling and Reenabling RARP Globally

To globally disable or reenable RARP on all router interfaces on which it is configured, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Reverse ARP**.	The Reverse ARP menu opens.
4. Choose **Globals**.	The Edit RARP Global Parameters window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **Enable/Disable** parameter. Click on **Help** or see the parameter description on page A-35.	
6. Click on **OK**.	You return to the Configuration Manager window.

Deleting RARP Globally

To globally delete RARP from all router interfaces on which it is configured, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Reverse ARP**.	The Reverse ARP menu opens.
4. Choose **Delete RARP**.	A confirmation window opens.
5. Click on **OK**.	Site Manager deletes RARP from all router interfaces and returns you to the Configuration Manager window.

Chapter 5
Customizing RIP Services

You customize the Routing Information Protocol (RIP) by setting RIP parameters as described under the following topics:

Topic	Page
Setting the RIP Diameter	5-2
Customizing a RIP Interface	5-3
Configuring RIP Accept and Announce Policies	5-24

Setting the RIP Diameter

The *RIP diameter* is a hop count that RIP uses to denote infinity. For RIP to operate properly, every router within the network must be configured with an identical RIP diameter value. If RIP is enabled, this parameter specifies the maximum number of hops within the autonomous system; if RIP is not enabled, IP uses the RIP diameter to determine network width.

You must set this parameter so that the interface cost, static cost, or route filter cost parameters do not exceed the RIP diameter.

The default RIP diameter value is 15 hops. Nortel Networks recommends that you accept the default value. However, you can use the BCC or Site Manager to specify a different RIP diameter value.

Using the BCC

To change the RIP diameter value, navigate to the global RIP prompt (for example, **box; ip; rip**) and enter:

rip-diameter *<infinity>*

infinity is a hop count indicating RIP infinity.

Using Site Manager

To change the RIP diameter value, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Global**.	The Edit IP Global Parameters window opens.
4. Set the **RIP Diameter** parameter. Click on **Help** or see the parameter description on page A-22.	
5. Click on **OK**.	You return to the Configuration Manager window.

Customizing a RIP Interface

When you add RIP to an IP interface, RIP is enabled with default values for all parameters. You customize RIP on the interface by modifying parameters as described in the following topics:

Topic	Page
Disabling and Reenabling RIP on an Interface	5-3
Selecting the RIP Version	5-4
Authenticating the Password on a Version 2 Update	5-6
Supplying RIP Updates on an Interface	5-8
Specifying the Update Mode	5-9
Sending Triggered Updates	5-11
Specifying a Time-to-Live Value	5-12
Receiving RIP Updates on an Interface	5-14
Supplying a Default Route on an Interface	5-15
Listening for a Default Route	5-17
Configuring a RIP Interface for Dial-Optimized Routing	5-18
Deleting RIP from an Interface	5-23

Disabling and Reenabling RIP on an Interface

You can use the BCC or Site Manager to disable or reenable RIP on the interface as required.

Using the BCC

To disable or reenable RIP on an interface, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

state <*state*>

state is one of the following:

enabled (default)
disabled

For example, the following command disables RIP on IP interface 2.2.2.2:

```
rip/2.2.2.2# state disabled
rip/2.2.2.2#
```

Using Site Manager

To disable or reenable RIP on an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to enable.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-64.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Selecting the RIP Version

You can specify whether RIP sends Version 1 updates, Version 2 updates with no aggregation of subnets, or Version 2 updates with subnet aggregation.

- In RIP Version 1 mode (the default), RIP generates Version 1 updates only, using the broadcast address as specified in RFC 1058. RIP aggregates subnet information.
- In RIP Version 2 mode, RIP generates Version 2 updates, using the multicast address 224.0.0.9 as specified in the RIP Version 2 RFC, 1388. RIP does not aggregate subnet information.
- In RIP Version 2 mode with aggregation, RIP generates Version 2 updates, using the multicast address, and performs aggregation of subnets into a natural network advertisement on interfaces belonging to another network.

In both Version 2 modes, RIP checks for a password on all received updates (see “Supplying a Default Route on an Interface” on page 5-15).

By default RIP sends Version 1 updates. You can use the BCC or Site Manager to choose Version 2 with or without aggregation of subnet information.

Using the BCC

To specify the RIP version, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

version <*version*>

version is one of the following:

rip1 (default)
rip2
aggr

For example, the following command causes RIP to send Version 2 updates with aggregation of subnets:

```
rip/2.2.2.2# version aggr
rip/2.2.2.2#
```

Using Site Manager

To specify the RIP version, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **RIP Mode** parameter. Click on **Help** or see the parameter description on page A-68.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Authenticating the Password on a Version 2 Update

By default, RIP running in Version 2 mode does not authenticate the password on incoming updates. RIP checks for the presence of a password as follows:

- If no password is present, RIP accepts the update.
- If a password is present, RIP drops the update.

With authentication enabled, RIP drops all received Version 1 updates and processes Version 2 updates in the following manner:

- If no password is present in the Version 2 update, RIP drops the update.
- If a password is present in the Version 2 update and that password is valid, RIP accepts the update.
- If the password is invalid, RIP drops the update.

You can use the BCC or Site Manager to configure a RIP interface for authentication and enable password checking. If you configure authentication on a RIP interface, you can assign the interface a 1- to 16-character password.

Using the BCC

To configure authentication for RIP Version 2, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

authentication-type *<type>*

type is one of the following:

none (default)
simple

To specify a password, enter:

authentication *<password>*

For example, the following command sequence enables authentication and sets the password to “secret” on IP interface 2.2.2.2:

```
rip/2.2.2.2# authentication-type simple
rip/2.2.2.2# authentication secret
```

Using Site Manager

To configure authentication for RIP Version 2, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the following parameters: • **Authentication Type** • **Authentication Password** Click on **Help** or see the parameter descriptions beginning on page A-69.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Supplying RIP Updates on an Interface

By default, RIP supplies RIP updates to neighboring networks on each interface. You can use the BCC or Site Manager to disable and reenable this feature on an interface as required.

Using the BCC

To disable or reenable RIP updates, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

supply *<state>*

state is one of the following:

enabled (default)
disabled

For example, the following command causes RIP to stop supplying updates on IP interface 2.2.2.2:

```
rip/2.2.2.2# supply disabled
rip/2.2.2.2#
```

Using Site Manager

To disable or reenable RIP updates, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **RIP Supply** parameter. Click on **Help** or see the parameter description on page A-64.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the Update Mode

RIP can issue routing updates in the following modes:

- Poisoned reverse (the default)
- Split horizon
- Actual cost

Poisoned reverse mode and *split horizon* mode are schemes for controlling the way a router advertises a route to the neighbor from which it learned the route.

In poisoned reverse updating, a router that sends updates to a neighbor includes routes learned from that neighbor but sets the route metric to infinity.

In split horizon updating, a router that sends updates to a neighbor omits routes that it learned from that neighbor.

On certain interfaces -- for example, on a frame relay interface that has virtual connections (VCs) to different routers that are part of the same logical IP subnet -- you may need to advertise all learned routes with the *actual cost*.

By default, RIP sends poisoned reverse updates. You can use the BCC or Site Manager to specify a different update mode.

Using the BCC

To specify the RIP update mode, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

mode *<mode>*

mode is one of the following:

poisoned
actual
split

For example, the following command causes RIP to send split-horizon updates on IP interface 2.2.2.2:

```
rip/2.2.2.2# mode split
rip/2.2.2.2#
```

Using Site Manager

To specify the RIP update mode, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Poisoned Reverse** parameter. Click on **Help** or see the parameter description on page A-66.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Sending Triggered Updates

RIP generates full routing updates at regular intervals. You can also configure RIP to generate an update on a specified interface each time it recalculates a route's metric. Such an update is called a *triggered update*. A triggered update contains only the routes that have changed. (RIP also sends full updates at regular intervals on interfaces configured for triggered updating.)

By default, triggered updates are disabled on all interfaces configured for RIP.

You can use the BCC or Site Manager to enable and disable triggered updates on this interface.

Using the BCC

To enable or disable triggered updates, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

triggered-updates *<state>*

state is one of the following:

enabled
disabled (default)

For example, the following command enables triggered updates on IP interface 2.2.2.2:

```
rip/2.2.2.2# triggered-updates enabled
rip/2.2.2.2#
```

Using Site Manager

To enable or disable triggered updates, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Triggered Updates** parameter. Click on **Help** or see the parameter description on page A-69.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Time-to-Live Value

By default, RIP inserts a time-to-live (TTL) value of one hop into each outbound routing update. Setting a TTL of 1 prevents RIP updates from inadvertently exiting the local network. Increasing the TTL introduces the risk of the update exiting the local network and being forwarded around the network.

Certain RIP implementations ignore packets with a TTL value of 1. Use this parameter to provide interoperability with such implementations.

You can use the BCC or Site Manager to specify a TTL value greater than 1.

Using the BCC

To specify a time-to-live value, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

ttl *<hops>*

hops is the time-to-live value (expressed as the number of hops) that RIP inserts in each outbound update.

For example, the following command causes RIP to insert a TTL value of 2 in each outbound update:

```
rip/2.2.2.2# ttl 2
rip/2.2.2.2#
```

Using Site Manager

To specify a time-to-live value, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Time to Live** parameter. Click on **Help** or see the parameter description on page A-67.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Receiving RIP Updates on an Interface

By default, RIP listens for routing updates on every interface on which it is enabled.

You can use the BCC or Site Manager to disable and reenable this feature as required.

Note: If you enable this feature, a configured policy can still prohibit the interface from updating its internal routing table.

Using the BCC

To specify whether RIP listens for routing updates, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

listen *<state>*

state is one of the following:

enabled (default)
disabled

For example, the following command causes RIP to stop listening for updates on IP interface 2.2.2.2:

```
rip/2.2.2.2# listen disabled
rip/2.2.2.2#
```

Using Site Manager

To specify whether RIP listens for routing updates, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear.
5. Set the **RIP Listen** parameter. Click on **Help** or see the parameter description on page A-65.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Supplying a Default Route on an Interface

When the routing table does not contain the route to a particular destination address, the router looks for a default route to the destination. Like any other route in the routing table, the default route can be acquired dynamically (by means of a routing protocol) or entered statically (by you).

This parameter is independent of the RIP supply parameter. A configured policy can override the parameter.

By default, RIP does not supply a default route. You can use the BCC or Site Manager to configure RIP to advertise an existing default route (one that is present in the routing table) in RIP updates sent to neighboring networks. You can also configure RIP to generate a default route if the routing table does not contain a default route.

Using the BCC

To specify whether an interface can supply or generate a default route, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

default-supply *<action>*

action is one of the following:

enabled
disabled (default)
generate

For example, the following command causes RIP to supply a default route on IP interface 2.2.2.2:

```
rip/2.2.2.2# default-supply enabled
rip/2.2.2.2#
```

Using Site Manager

To specify whether an interface can supply or generate a default route, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Default Route Supply** parameter. Click on **Help** or see the parameter description on page A-65.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Listening for a Default Route

By default, RIP ignores inbound advertisements of a default route (0.0.0.0) on the interfaces where it is configured. You can use the BCC or Site Manager to configure RIP to listen for a default route and add that information to its internal routing table.

Using the BCC

To specify whether RIP listens for the default route, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

default-listen *<action>*

action is one of the following:

enabled
disabled (default)

For example, the following command causes RIP to listen for a default route on IP interface 2.2.2.2:

```
rip/2.2.2.2# default-listen enabled
rip/2.2.2.2#
```

Using Site Manager

To specify whether RIP listens for the default route, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **Default Route Listen** parameter. Click on **Help** or see the parameter description on page A-66.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a RIP Interface for Dial-Optimized Routing

Dial-optimized routing is a method for reducing costs on dialed lines. Under dial-optimized routing, RIP exchanges routing information on the interface only when the router or a peer has activated the connection for a data transmission. RIP does not initiate a connection on a dialed line solely for the purpose of issuing a routing update.

For information about enabling dial-optimized routing, see *Configuring Dial Services*.

After you enable dial-optimized routing, you can set RIP timers to control the way RIP generates updates on interfaces to dialed lines. Setting RIP timers allows you to specify the following:

- The frequency at which RIP broadcasts full RIP updates on the interface
- The timeout period that RIP waits before considering a network unreachable
- The holddown period that unreachable routes are retained in the routing table and advertised

Specifying an Update Interval

By default, RIP generates a full update every 30 seconds on each interface configured with RIP. If you enabled dial-optimized routing on this interface, the default is 1 hour.

You can use the BCC or Site Manager to specify an update interval for the interface. If you change the value of the update interval from the default value, you should also change the timeout period to a value equal to three times the update interval (see "Specifying a Timeout Period" on page 5-20 and "Specifying a Holddown Period" on page 5-21).

Using the BCC

To specify the frequency at which RIP generates a full update, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

broadcast-timer *<seconds>*

seconds is the broadcast interval in seconds. The default interval is 30 seconds.

For example, the following command causes RIP to broadcast a full update every 15 seconds on IP interface 2.2.2.2:

```
rip/2.2.2.2# broadcast-timer 15
rip/2.2.2.2#
```

Using Site Manager

To specify the frequency at which RIP generates a full update, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Broadcast Timer** parameter. Click on **Help** or see the parameter description on page A-67.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Timeout Period

By default, RIP waits 90 seconds for an update from a network before it considers that network to be unreachable. If you enabled dial-optimized routing on this interface, the default is 3 hours.

You can use the BCC or Site Manager to specify a timeout period from 15 seconds through 259,200 seconds (72 hours). For dial-optimized routing the maximum value is 3,628,800 seconds (6 weeks).

You should set the timeout timer to a value equal to three times the value of the update interval (see "Specifying an Update Interval" on page 5-18).

Using the BCC

To specify the timeout period, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

timeout-timer <*seconds*>

seconds is the timeout interval in seconds.

For example, the following command sets the RIP timeout timer to 120 seconds on IP interface 2.2.2.2:

```
rip/2.2.2.2# timeout-timer 120
rip/2.2.2.2#
```

Using Site Manager

To specify the timeout period, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Timeout Timer** parameter. Click on **Help** or see the parameter description on page A-67.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Holddown Period

When RIP determines that a network is unreachable, RIP continues to advertise a route to that network for a default holddown period of 90 seconds. If you enabled dial-optimized routing on this interface, the default is 3 hours.

You can use the BCC or Site Manager to specify a holddown period from 15 seconds through 259,200 seconds (72 hours). For dial-optimized routing, the maximum value is 3,628,800 seconds (6 weeks).

You should set the holddown timer to a value equal to three times the value of the update interval (see "Specifying an Update Interval" on page 5-18).

Using the BCC

To specify the holddown period, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

holddown-timer *<seconds>*

seconds is the holddown period expressed in seconds.

For example, the following command sets the holddown timer to 60 seconds on IP interface 2.2.2.2:

```
rip/2.2.2.2# holddown-timer 60
rip/2.2.2.2#
```

Using Site Manager

To specify the holddown period, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Holddown Timer** parameter. Click on **Help** or see the parameter description on page A-68.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Stabilization Time

The stabilization time is the period that RIP allows itself to learn all routes from its neighbors before sending full updates. By default, RIP uses a stabilization time of 120 seconds.

You can use Site Manager to specify a stabilization time for an interface.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RIP Interfaces**.	The IP RIP Interface Configuration window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Click on the RIP interface that you want to edit.	The parameter values for that interface appear in the IP RIP Interface Configuration window.
5. Set the **Initial Stabilization Timer** parameter. Click on **Help** or see the parameter description on page A-70.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Deleting RIP from an Interface

To delete RIP from an interface on which it is currently configured, proceed as follows.

Using the BCC

To delete RIP from an interface, navigate to a RIP interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; rip**) and enter:

delete

For example, the following command deletes RIP from interface 2.2.2.2:

```
rip/2.2.2.2# delete
ip/2.2.2.2/255.255.255.0#
```

Using Site Manager

To delete RIP from an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on the connector from which you want to delete RIP services.	The Edit Connector window opens.
2. Click on **Edit Circuit**.	The Circuit Definition window opens.
3. Choose **Protocols**.	The Protocols menu opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Choose **Add/Delete**.	The Select Protocols window opens. The RIP button is checked to show that RIP is enabled on the circuit.
5. Click on **RIP**.	
6. Click on **OK**.	Site Manager deletes RIP services from the connector and returns you to the Circuit Definition window.
7. Choose **File**.	The File menu opens.
8. Choose **Exit**.	You return to the Configuration Manager window.

Configuring RIP Accept and Announce Policies

RIP accept policies and announce policies allow you to control the flow of routing information in and out of the routing table as follows:

- An *accept policy* controls the routing information that is considered for inclusion in the IP routing table.
- An *announce policy* controls the routing information that RIP advertises.

For an introduction to IP policies, see “IP Routing Policies and Filters” on page 1-22.

The following topics describe how to configure RIP accept and announce policies:

Topic	Page
Defining a RIP Accept Policy	5-25
Supplying Modification Values for a RIP Accept Policy	5-28
Specifying Matching Criteria for a RIP Accept Policy	5-29
Defining a RIP Announce Policy	5-31
Supplying Modification Values for a RIP Announce Policy	5-33
Specifying Matching Criteria for a RIP Announce Policy	5-35

Defining a RIP Accept Policy

To define a new RIP accept policy, you must do the following:

- Supply a name for the policy.
- Specify whether RIP accepts or ignores an update that matches the policy.
- Rank the policy according to preference, precedence, and other criteria.

You can use the BCC or Site Manager to define a RIP accept policy.

Using the BCC

To define a new policy, navigate to the RIP global prompt (for example, **box; ip; rip**) and enter:

accept *<policy_name>*

policy_name is a unique name for the RIP accept policy.

A policy-specific prompt appears, indicating that the BCC has created the policy using default values for all parameters.

For example, the following command creates an accept policy named pol_1:

```
rip# accept pol_1
accept/pol_1/rip#
```

To customize the default values for a policy, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 5-1.

Table 5-1. BCC Definition Parameters for RIP Accept Policies

Parameter	Values	Function
state	enabled (default) disabled	Enables and disables the policy that you created
action	ignore (default) accept	Specifies whether the protocol ignores a route that matches the policy or forwards the route to the routing table manager
preference	1 (default) to 16	Assigns a metric value (the higher the number, the greater the preference) to a route that the protocol forwards to the routing table manager. If confronted with multiple routes to the same destination, the routing table manager may use this value to decide which route to insert. Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference, and routes for the most specific networks (longest address and mask) should have the highest preference.
precedence	0 (default) to any integer	Assigns a metric value to this policy (a policy with a higher value takes precedence over a policy with a lower value). This value determines the order of precedence for policies that match the same route.

For example, the following command sets the state to disabled for RIP accept policy pol_1:

```
accept/pol_1/rip# state disabled
accept/pol_1/rip#
```

Using Site Manager

To create an accept policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters**.	The IP policies menu opens.
4. Choose **RIP.**	The RIP policies menu opens.
5. Choose **Accept Policies**.	The RIP Accept Policy Filters window opens.
6. Click on **Add**.	The RIP Accept IP Policy Filter Configuration window opens.
7. Set the following parameters: • **Name** • **Action** • **Route Preference** • **Rule Precedence** • **Networks** • **From Gateway** • **Received on Interface** • **Apply Subnet Mask** Click on **Help** or see the parameter descriptions beginning on page B-2.	
8. Click on **OK**.	You return to the RIP Accept Policy Filters window.
9. Click on **Apply**, and then on **Done**.	You return to the Configuration Manager window.

Supplying Modification Values for a RIP Accept Policy

You can use the BCC or Site Manager to supply values that RIP uses to modify fields in a RIP update that matches the policy.

Using the BCC

To supply a modification value to an update that matches a policy, navigate to the policy prompt (for example, **box; ip; rip; accept/pol_1**) and enter:

modify

A modification prompt appears for the policy.

To specify a value, enter:

<parameter> <value>

parameter value is the parameter/value pair shown in Table 5-2.

Table 5-2. BCC Modification Parameter for RIP Accept Policies

Parameter	Values	Function
mask	0.0.0.0 or an IP mask	Specifies a mask that will override the interface's subnet mask in the presence of networks with variable-length subnet masks

For example, the following command sequence specifies an override mask of 255.0.0.0 for accept policy pol_1:

```
accept/pol_1/rip# modify
modify/rip/accept/pol_1# mask 255.0.0.0
modify/rip/accept/pol_1#
```

Using Site Manager

To specify a modification value to an update that matches a policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters**.	The IP policies menu opens.
4. Choose **RIP.**	The RIP policies menu opens.
5. Choose **Accept Policies**.	The RIP Accept Policy Filters window opens.
6. Select the policy that you want to edit.	
7. Set the **Apply Subnet Mask** parameter. Click on **Help** or see the parameter description on page B-6.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying Matching Criteria for a RIP Accept Policy

You can use the BCC or Site Manager to specify matching criteria for a policy.

Using the BCC

To specify matching criteria for an accept policy, navigate to the policy prompt (for example, **box; ip; rip; accept/pol_1**) and enter:

match

A match prompt appears for the policy.

To specify matching criteria, enter:

<match_criterion> <ip_address>

match_criterion is one of the BCC match objects listed in Table 5-3.

Table 5-3. BCC Match Parameters for RIP Accept Policies

Object	Meaning
network*	Specifies the networks that match this policy. Each specification consists of a network IP address, a mask, and a flag to indicate whether the network address refers to a specific network or to a range of networks. Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Use the default empty list to match any route.
rip-gateway*	Specifies the addresses of one or more routers that could send RIP updates to this router. This policy matches RIP advertisements from routers on this list. Use the default empty list to indicate that this policy applies to RIP updates from any router.
rip-interface*	Specifies the IP addresses of one or more interfaces on this router. This policy matches RIP updates received on interfaces that appear in this list. Use the default empty list to indicate that this policy applies to RIP updates received on any interface.

* The BCC implements this matching criterion as an object. To specify a list, create multiple objects.

For example, the following command specifies RIP interface 2.2.3.3 on this router as a match for RIP accept policy pol_1:

```
accept/pol_1/rip# match
match/rip/accept/pol_1# rip-interface 2.2.3.3
rip-interface/2.2.3.3/rip/accept/pol_1#
```

Using Site Manager

To specify matching criteria for an accept policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters**.	The IP policies menu opens.
4. Choose **RIP.**	The RIP policies menu opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Choose **Accept Policies**.	The RIP Accept Policy Filters window opens.
6. Select the policy that you want to edit.	
7. Set the following parameters: • **Networks** • **From Gateway** • **Received on Interface** Click on **Help** or see the parameter descriptions on pages B-3 and B-5.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Defining a RIP Announce Policy

To define a new RIP announce policy, you must do the following:

- Supply a name for the policy.
- Specify whether RIP announces or ignores an update that matches the policy.
- Rank the policy according to precedence and other criteria.

You can use the BCC or Site Manager to define a RIP announce policy.

Using the BCC

To define a new announce policy, navigate to the RIP global prompt (for example, **box; ip; rip**) and enter:

announce *<policy_name>*

policy_name is a unique name for the announce policy.

A policy-specific prompt appears, indicating that the BCC has created the policy, using default values for all parameters.

For example, the following command creates a RIP announce policy named pol_1:

```
rip# announce pol_1
announce/pol_1/rip#
```

To customize the default values for the policy, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 5-4.

Table 5-4. BCC Definition Parameters for RIP Announce Policies

Parameter	Values	Function
state	enabled (default) disabled	Enables or disables this policy
action	ignore (default) announce	Specifies whether or not to advertise a route that matches this policy
precedence	0 (default) to any metric value	Specifies a metric value to compare this policy with other policies that a route may match. A policy with a higher metric takes precedence over a policy with a lower metric. In case of a tie, the protocol uses an internal index value assigned to the policy by IP software.

For example, the following command specifies a precedence value of 12 for RIP announce policy pol_1:

```
announce/pol_1/rip# precedence 12
announce/pol_1/rip#
```

Using Site Manager

To create an announce policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters**.	The IP policies menu opens.
4. Choose **RIP.**	The RIP policies menu opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Choose **Announce Policies**.	The RIP Announce Policy Filters window opens.
6. Click on **Add**.	The RIP Announce IP Policy Filter Configuration window opens.
7. Set the following parameters: • **Name** • **Action** • **Rule Precedence** Click on **Help** or see the parameter descriptions beginning on page B-8.	
8. Click on **OK**.	You return to the RIP Announce Policy Filters window.
9. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Supplying Modification Values for a RIP Announce Policy

You can use the BCC or Site Manager to supply a value that RIP uses to modify fields in a RIP update that matches the policy.

Using the BCC

To supply a modification value to an update that matches a policy, navigate to the policy prompt (for example, **box; ip; rip; announce/pol_1**) and enter:

modify

A modification prompt appears for the policy.

To specify an override value, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 5-5.

Table 5-5. BCC Modification Parameters for RIP Announce Policies

Parameter	Values	Function
metric	0 (the default) or an export metric	Specifies an optional RIP export metric to use when advertising a route that matches this policy
advertise*	List of network identifiers	Specifies network IDs to include in place of the network IDs listed in the route to be advertised

* The BCC implements this criterion as an object. To specify a list, create multiple objects.

For example, the following command sequence specifies an override metric of 5 for routes that match announce policy pol_1:

```
announce/pol_1/rip# modify
modify/rip/announce/pol_1# metric 5
```

Using Site Manager

To specify an override value for an announce policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters**.	The IP policies menu opens.
4. Choose **RIP.**	The RIP policies menu opens.
5. Choose **Announce Policies**.	The RIP Announce Policy Filters window opens.
6. Select the policy that you want to edit.	
7. Set the following parameters: • **Advertise** • **RIP Metric** Click on **Help** or see the parameter descriptions beginning on page B-11.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying Matching Criteria for a RIP Announce Policy

You can use the BCC or Site Manager to specify matching criteria for a RIP announce policy.

Using the BCC

To specify matching criteria for a RIP announce policy, navigate to the policy prompt (for example, **box; ip; rip; announce/pol_1**) and enter:

match

A match prompt for the policy appears.

For example, the following command invokes a match prompt for the RIP announce policy pol_1:

```
announce/pol_1/rip# match
match/rip/announce/pol_1#
```

At this prompt, enter:

<match_criterion>

match_criterion is one of the parameter/value pairs listed in Table 5-6.

Table 5-6. BCC Match Parameters for RIP Announce Policies

Parameter	Values	Function
external-source	any (default) direct static rip ospf egp bgp	Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy. This parameter applies only to OSPF routes that use the new ASE type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.

(continued)

Table 5-6. BCC Match Parameters for RIP Announce Policies *(continued)*

Parameter	Values	Function
ospf-type	any (default) type1 type2 external internal	Specifies which types of OSPF routes match this policy, and applies only to OSPF-sourced routes when OSPF is included as a route source
protocol-source	any (default) direct static rip ospf egp bgp	Specifies one or more route source identifiers. If you select a route source ID, a route from that source that meets the other criteria of this policy matches the policy.
bgp-as*	List of AS numbers	Specifies one or more autonomous system numbers. This policy applies to BGP advertisements received from BGP peers in an AS on this list, and applies only to BGP-sourced routes when BGP is included as a route source.
bgp-next-hop*	List of IP addresses	Specifies one or more IP addresses. This policy applies to BGP advertisements whose Next Hop attribute matches an IP address on this list, and applies only to BGP-sourced routes when BGP is included as a route source.
bgp-peer*	List of IP addresses	Specifies the IP address of one or more BGP peers. This policy applies to BGP advertisements authored by a router on this list, and applies only to BGP-sourced routes when BGP is included as a route source.
egp-as*	List of AS numbers	Specifies one or more autonomous system numbers. This policy applies to EGP advertisements received from EGP peers in an AS on this list, and applies only to EGP-sourced routes when EGP is included as a route source.

(continued)

Table 5-6. BCC Match Parameters for RIP Announce Policies *(continued)*

Parameter	Values	Function
egp-gateway*	List of IP addresses	Specifies the IP address of one or more EGP gateways. This policy applies to EGP advertisements that use a gateway on this list as the next hop, and applies only to EGP-sourced routes when EGP is included as a route source.
egp-peer*	List of IP addresses	Specifies the IP address of one or more EGP peers. This policy applies to EGP advertisements authored by a router on this list, and applies only to EGP source routes when EGP is included as a route source.
inbound-interface*	List of IP addresses	Specifies the address of one or more interfaces on this router. This policy applies to RIP advertisements received on the interfaces on this list, and applies only to RIP-sourced routes when RIP is included as a route source.
network*	List of IP addresses	Specifies which networks will match this policy. Each identifier consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or to a range of networks. Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Enter an empty list to match any route.
ospf-router-id*	List of IP addresses	Specifies the IDs of one or more OSPF routers. This policy applies to OSPF advertisements authored by a router on this list, and applies only to OSPF-sourced routes when OSPF is included as a route source.
ospf-tag*	List of OSPF ASE tag values	Specifies tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain tag values in this list, and applies only to OSPF-sourced ASE routes when OSPF is included as a route source.

(continued)

Table 5-6. BCC Match Parameters for RIP Announce Policies *(continued)*

Parameter	Values	Function
outbound-interface*	List of IP addresses	Specifies a list of outbound RIP interfaces. If an interface appears on this list, the policy applies to RIP advertisements sent via that interface.
rip-gateway*	List of IP addresses	Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list, and applies only to RIP-sourced routes when RIP is included as a route source.

* The BCC implements this matching criterion as an object. To specify multiple items, create multiple objects.

Using Site Manager

To specify matching criteria for a RIP announce policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters**.	The IP policies menu opens.
4. Choose **RIP.**	The RIP policies menu opens.
5. Choose **Announce Policies**.	The RIP Announce Policy Filters window opens.
6. Select the policy that you want to edit.	

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
7. Set the following parameters: • **Networks** • **Route Source** • **External Route Source** • **Advertise** • **From RIP Gateway** • **Received on RIP Interface** • **Outbound Interfaces** • **From OSPF Router ID** • **Received OSPF Type** • **Received OSPF Tag** • **From EGP Peer** • **From EGP AS** • **Received EGP Gateway** • **From BGP Peer** • **From BGP AS** • **Received BGP Next Hop** Click on **Help** or see the parameter descriptions beginning on page B-8.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Chapter 6
Customizing OSPF Services

You customize OSPF by setting OSPF parameters as described under the following topics:

Topic	Page
Customizing OSPF Global Parameters	6-2
Customizing OSPF on an IP Interface	6-18
Defining an Area	6-40
Configuring an NSSA Address Range	6-50
Configuring an Area Border Router	6-53
Configuring OSPF Accept and Announce Policies	6-57

Customizing OSPF Global Parameters

OSPF global parameters affect the way OSPF runs on the router. They apply to all OSPF interfaces. You customize OSPF global parameters as described under the following topics:

Topic	Page
Enabling and Disabling OSPF on the Router	6-3
Supplying an OSPF Router ID	6-4
Configuring the Soloist and Backup Soloist on a Slot	6-5
Enabling the Boundary Function	6-7
Configuring the Metric Type for an ASE Advertisement	6-8
Choosing a Tag Generation Method for an ASE Advertisement	6-11
Setting the Holddown Timer	6-14
Configuring Message Logging	6-15
Configuring External Route Preference	6-17

Enabling and Disabling OSPF on the Router

When you start OSPF on the router, OSPF is automatically enabled.

You can use the BCC or Site Manager to disable and reenable OSPF on the router.

Using the BCC

To disable or reenable OSPF on the router, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

state *<state>*

state is one of the following:

enabled (default)
disabled

Using Site Manager

To disable or reenable OSPF on the router, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-37.	The value you chose appears in the Enable field.
6. Click on **OK**.	You return to the Configuration Manager window.

Supplying an OSPF Router ID

Each router configured with OSPF has an OSPF router ID. This IP address uniquely identifies this router in the OSPF domain.

By convention, and to ensure uniqueness, the router ID should be one of the router's IP interface addresses.

The router ID determines the designated router on a broadcast link if the priority values of the routers being considered are equal. The higher the router ID, the greater its priority.

If both OSPF and BGP are running on the router, the OSPF router ID must be identical to the BGP identifier. In addition, the OSPF router ID must match one of the IP addresses configured on the router.

By default, OSPF uses the IP address of the first OSPF circuit configured on this router. You can use the BCC or Site Manager to specify an IP address.

Using the BCC

To change the OSPF router ID, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

router-id <*ip_address*>

ip_address is a valid IP address in dotted-decimal notation.

Using Site Manager

To change the OSPF router ID, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **Router ID** parameter. Click on **Help** or see the parameter description on page A-38.	
6. Click on **OK**.	You return to the Configuration Manager window.

Configuring the Soloist and Backup Soloist on a Slot

The OSPF protocol is implemented as a *soloist* -- that is, as a single process running on a single slot of a router. When you add an OSPF interface to a circuit, the router enables OSPF on a slot. If the slot on which the OSPF soloist is running goes down, the router tries to run OSPF on another slot.

Each time the OSPF soloist is restarted, all of the routing information is lost and must be relearned from the network. The *OSPF backup soloist* provides a method of preserving information learned from the network in the event of an OSPF crash or slot removal, avoiding the time-consuming and resource-intensive process of relearning routing information. In the event of a crash or slot removal, transition between the OSPF primary and backup soloists occurs without relearning routing information from the network.

By default, the router uses any available slot for the OSPF soloist. You can use the BCC or Site Manager to specify a slot.

By default, OSPF does not maintain a copy of the link state database (LSDB) for the backup soloist. You can use the BCC or Site Manager to enable this feature.

Using the BCC

To specify a slot for the OSPF soloist, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

slot-mask *<slot>*

slot is **all-slots** (the default) or an integer from 1 through 14 to indicate a slot.

To maintain a copy of the LSDB for the backup soloist, navigate to the global OSPF prompt and enter:

backup-lsdb enabled

For example, the following command sequence specifies slot 12 for the OSPF soloist and causes OSPF to maintain a separate copy of the LSDB for the backup soloist:

```
ospf# slot-mask 12
ospf# backup-lsdb enabled
ospf#
```

Using Site Manager

To specify a slot for the soloist and maintain a copy of the LSDB for the backup soloist, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the following parameters: • **OSPF Slot** • **Backup Enable** Click on **Help** or see the parameter descriptions beginning on page A-39.	
6. Click on **OK**.	You return to the Configuration Manager window.

Enabling the Boundary Function

An OSPF boundary router does the following:

1. Receives information about routes outside the OSPF AS (using BGP, RIP, or another routing protocol)
2. Formats this information in AS external (ASE) advertisements
3. Propagates the ASEs into the OSPF domain (using OSPF)

By default, the boundary router function is disabled on the router. You can use the BCC or Site Manager to configure the router as an OSPF boundary router.

Using the BCC

To configure the router as a boundary router, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

as-boundary-router *<state>*

state is one of the following:

true
false (default)

Using Site Manager

To configure the router as a boundary router, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the **AS Boundary Router** parameter. Click on **Help** or see the parameter description on page A-38.	
6. Click on **OK**.	You return to the Configuration Manager window.

Configuring the Metric Type for an ASE Advertisement

Note: This parameter applies to boundary routers only.

Each ASE that a boundary router injects into the AS includes a type 1 or type 2 metric. The type 1 metric is equivalent to the metric of the non-OSPF route. The type 2 metric is either the metric of the non-OSPF route or the weight value calculated for that route.

Figure 6-1, for example, shows three routers in an OSPF domain. Router A and router B are both configured to generate ASE routes using the route weight as the type 2 metric. Router A and router B both learn a route to destination X. The following steps occur:

1. Boundary router A learns a route to destination X via EGP.
2. Boundary router A advertises the route to the internal router as an OSPF ASE route. The type 2 metric in the advertisement contains the route weight value calculated for the EGP route to destination X.
3. Boundary router B learns a route to destination X via BGP.
4. Boundary router B advertises the route to the internal router as an OSPF ASE route. The type 2 metric in the advertisement contains the route weight value calculated for a BGP route.
5. To determine the preferred route, the internal router compares the type 2 metrics -- the EGP route weight and the BGP route weight.
6. The internal router chooses the BGP route -- the route with the lower weight.

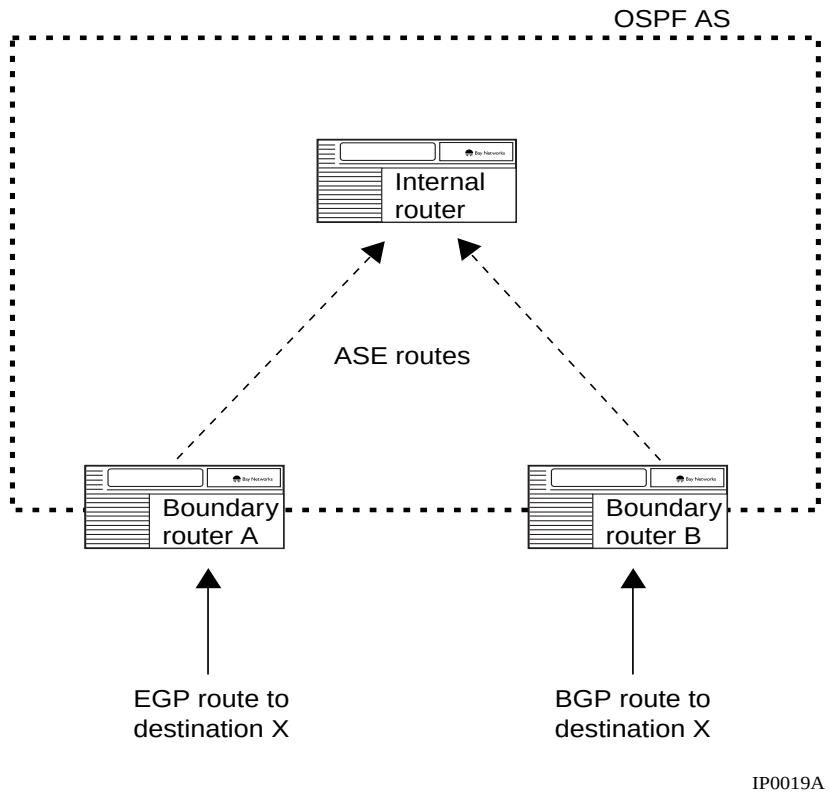

Figure 6-1. OSPF ASE Routes

By default, an OSPF boundary router generates a type 2 metric for BGP, EGP, and RIP routes. For routes from all other sources, the boundary router generates a type 1 metric.

Note: The route weight value will appear to be greater than the route's original metric. For this reason, all routers advertising a particular network must use the same metric type -- type 1 or type 2. If not, the router that receives the advertisements may choose the wrong route.

Using the BCC or Site Manager, you can configure a boundary router to use the route weight as the OSPF metric.

Using the BCC

To enable or disable the use of the route weight as the OSPF metric, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

ase-metric-support *<state>*

state is one of the following:

enabled
disabled (default)

Using Site Manager

To enable or disable the use of the route weight as the OSPF metric, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **ASE Metric Support** parameter. Click on **Help** or see the parameter description on page A-39.	
6. Click on **OK**.	You return to the Configuration Manager window.

Choosing a Tag Generation Method for an ASE Advertisement

Note: This parameter applies to boundary routers only.

An OSPF AS external route advertisement includes an external route tag field. This field allows boundary routers in an AS to exchange information about external routes. (The specific nature of this information is outside the scope of OSPF.)

By default, Nortel Networks boundary routers that generate ASE advertisements set the external route tag field to 0. For a boundary router running OSPF and BGP, you can configure OSPF to set the external route tag field with a value in accordance with RFC 1403, *OSPF/BGP Interaction.*

In Figure 6-2, for example, boundary router R14 running OSPF and BGP learns external routes via BGP:

1. Router R14 generates an ASE advertisement describing the route. OSPF fills in the external route tag with BGP-specific information according to RFC 1403.
2. Router R14 injects the ASE into the AS, and OSPF routers flood the ASE throughout the AS.
3. Router R1, which runs OSPF and BGP, receives the ASE. R1 generates a BGP update, using the contents of the external route tag to set the Origin and AS Path attributes in the update.

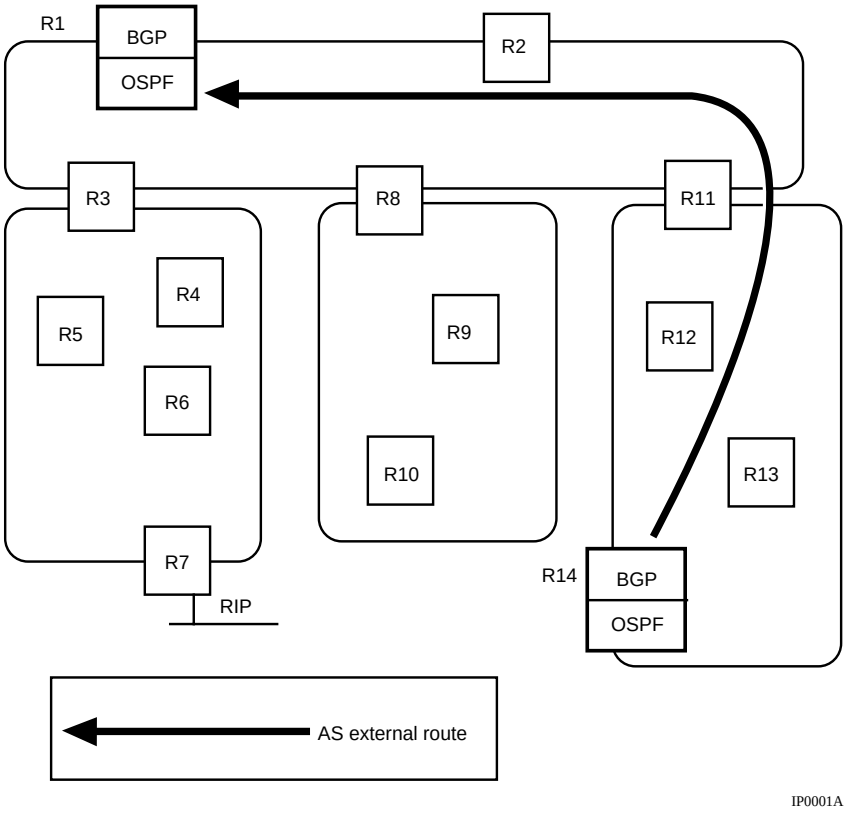

Figure 6-2. AS External Route Tag

You can use the BCC or Site Manager to choose the tag generation method.

Using the BCC

To set the tag generation method for ASE advertisements, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

as-default-tag *<method>*

method is one of the following:

default (default)
automatic
proprietary (reserved for debugging purposes)

Using Site Manager

To set the tag generation method for ASE advertisements, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Tag Generation Method** parameter. Click on **Help** or see the parameter description on page A-41.	
6. Click on **OK**.	You return to the Configuration Manager window.

Setting the Holddown Timer

The holddown timer controls how often OSPF calculates a route. Its purpose is to free up the CPU. A value of 0 means there is no holddown time.

By default, the holddown timer is set to 1 second. You can use the BCC or Site Manager to specify a different value.

Using the BCC

To set the holddown timer value, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

holddown *<value>*

value is 0 (no holddown time) or the holddown time in seconds.

For example, the following command specifies 5 seconds as the OSPF holddown time:

```
ospf# holddown 5
ospf#
```

Using Site Manager

To set the holddown timer value, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Hold Down Timer** parameter. Click on **Help** or see the parameter description on page A-39.	
6. Click on **OK**.	You return to the Configuration Manager window.

Configuring Message Logging

OSPF maintains a primary log and a backup log for OSPF messages.

By default, OSPF logs the following messages in the primary log:

- Trace
- Info
- Debug
- INTF State
- NBR State
- Bad LS

By default, OSPF logs no messages in the backup log.

You can use the BCC or Site Manager to specify the messages that OSPF writes to the primary and backup logs.

OSPF can log any or all of the messages described in Table 6-1.

Table 6-1. OSPF Log Messages

Message	Example
Trace	Designated Router changed on network: x.x.x.x x.x.x.x -> x.x.x.x
Info	OSPF enabled
OSPF debug	OSPF couldn't get a buffer, dying
INTF state	Interface x.x.x.x up on circuit x
NBR state	T2: Neighbor x.x.x.x Event: x State change: x->x
LSA self-origin	T4: Originating new LSA - type x LSID x.x.x.x router x.x.x.x
LSA receipt	T5: Received new LSA - type x LSID x.x.x.x router x.x.x.x neighbor x.x.x.x
Route change	T6: Routing Table changed - type x destination x.x.x.x old x.x.x.x new x.x.x.x
Bad LS	R4: Ack received for non-existent LSA: type x LSID x.x.x.x neighbor x.x.x.x

(continued)

Table 6-1. OSPF Log Messages *(continued)*

Message	Example
Less recent LSA	C3: Packet Rejected: LS UPDATE: LESS RECENT RX (x) src x.x.x.x type x ls_id: x.x.x.x adv_rtr: x.x.x.x ls_seq: x ls_age: x db_seq: x db_age: x elapse: x freeme:x ackcnt:x nbr_retrans:x nbrEcnt:x Fcnt:x
More recent LSA	R3: Received more recent self-originated LSA: type x LSID x.x.x.x router x.x.x.x neighbor x.x.x.x
Max age LSA	N3: LSA of MaxAge flushed: type x LSID x.x.x.x router x.x.x.x

Using the BCC

To specify the messages that OSPF writes to the primary log, navigate to the global OSPF prompt (for example, **box; ip; ospf**) and enter:

log-mask *<mask>*

To specify the messages that OSPF writes to the backup log, navigate to the global OSPF prompt and enter:

backup-log-mask *<mask>*

mask is a bit sequence indicating the messages that you want to log.

Using Site Manager

To specify the messages that OSPF writes to the primary log, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
5. Set the following parameters: • **Primary Log Mask** • **Backup Log Mask** Click on **Help** or see the parameter descriptions beginning on page A-40.	
6. Click on **OK**.	You return to the Configuration Manager window.

Configuring External Route Preference

By default, when OSPF receives multiple AS external LSAs for the same destination, OSPF applies the preference rules specified by RFC 1583 (superseded by RFC 2178).

You can use Site Manager to configure OSPF to apply the preference rules specified by RFC 2178. These rules are designed to prevent routing loops when AS external LSAs for the same destination originate from different areas.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **RFC 1583 Compatibility Enable** parameter. Click on **Help** or see the parameter description on page A-42.	
6. Click on **OK**.	You return to the Configuration Manager window.

Customizing OSPF on an IP Interface

When you add OSPF to an IP interface, OSPF is configured on the interface with default values for all interface parameters. You customize OSPF on the interface by modifying values as described under the following topics:

Topic	Page
Enabling and Disabling OSPF on an Interface	6-19
Configuring an Area ID	6-20
Specifying the Interface Type	6-21
Specifying Router Priority for a Multiaccess Network	6-23
Estimating the Transit Delay	6-24
Setting the Retransmit Interval	6-25
Setting the Hello Interval	6-27
Setting the Dead Interval	6-29
Setting the Poll Interval for NBMA Neighbors	6-30
Specifying the Metric Cost	6-31
Specifying the MTU Size	6-34
Specifying the MTU Mismatch Detect	6-35
Configuring a Neighbor on an NBMA Interface	6-36
Deleting OSPF from an IP Interface	6-38

Enabling and Disabling OSPF on an Interface

When you add OSPF to an IP interface, OSPF is automatically enabled on that interface. The interface will be advertised as an internal route. In addition, the interface can be used to form a neighbor relationship.

You can use the BCC or Site Manager to disable and reenable OSPF on an interface.

Using the BCC

To disable or reenable OSPF on an IP interface, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

state <*state*>

state is one of the following:

enabled (default)
disabled

Using Site Manager

To disable or reenable OSPF on an IP interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to modify.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-42.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring an Area ID

In an AS that contains multiple areas, each OSPF interface is configured with the ID of the area to which it is connected.

You supply an area ID when you add OSPF to the interface. You can use the BCC or Site Manager to assign a different area ID to the interface.

Using the BCC

To change an area ID, navigate to the OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

area <*area_id*>

area_id is an area identifier in dotted-decimal format.

For example, the following command assigns area ID 0.0.0.2 to interface 2.2.2.2. (This means that interface 2.2.2.2 connects the router to a network in OSPF area 0.0.0.2.)

```
ospf/2.2.2.2# area 0.0.0.2
ospf/2.2.2.2#
```

Using Site Manager

To change an area ID for an OSPF interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
6. Set the **Area ID** parameter. Click on **Help** or see the parameter description on page A-43.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the Interface Type

OSPF supports the following types of interfaces:

- *Broadcast.* An interface connected to a broadcast network that supports multiple routers. OSPF can address a single physical message to all attached routers. Examples of such a network are Ethernet, FDDI, and token ring.
- *Nonbroadcast multi-access* (NBMA). An interface to an NBMA network that supports multiple routers. However, OSPF cannot address a single physical message to all routers. Examples of such a network are frame relay and X.25.
- *Point-to-point.* An interface to a point-to-point network that joins a single pair of OSPF routers. An example of such a network is a network of synchronous lines or a direct mode frame relay environment.
- *Point-to-multipoint.* An interface to a point-to-multipoint network that supports multiple routers in a partial mesh configuration. An example of such a network is a group mode frame relay environment. Nortel Networks supports the standard OSPF point-to-multipoint interface and also provides a proprietary point-to-multipoint solution.
- *Passive*. A passive interface only receives advertisements. OSPF cannot use it to form neighbor relationships, accept Hello messages, or send advertisements. On other interfaces, OSPF advertises the network attached to a passive interface as a stub network.

Note: If the interface is connected to an NBMA network, you need to configure neighbors manually. For instructions on manually configuring neighbors, see “Configuring a Neighbor on an NBMA Interface” on page 6-36.

By default, OSPF assumes that the interface is attached to a broadcast network. You can use the BCC or Site Manager to specify another network type.

Using the BCC

To change the interface type, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

type <*interface_type*>

interface_type is one of the interface types described in Table 6-2.

Table 6-2. OSPF Interface Types

Type	Description
broadcast	Default. Choose broadcast if this network is a broadcast LAN, such as Ethernet.
nbma	Choose nbma (nonbroadcast multi-access) if the network is a nonbroadcast network, such as X.25.
pointopoint	Choose point-to-point for a synchronous, point-to-point interface.
ietf	Choose ietf if the network is a point-to-multipoint network. It is recommended that you use this network type rather than the Nortel Networks proprietary point-to-multipoint network.
pmp	Choose pmp (point-to-multipoint) to use the Nortel Networks proprietary point-to-multipoint solution.
passive	Choose passive to configure an interface that OSPF cannot use to form neighbor relationships. OSPF cannot accept Hello messages or send advertisements on the passive interface.

Using Site Manager

To change the interface type, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Type** parameter. Click on **Help** or see the parameter description on page A-43.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying Router Priority for a Multiaccess Network

The router priority value is used in multiaccess networks (broadcast, NBMA, or point-to-multipoint) to elect the designated router. A router with a priority of 0 is not eligible to become the designated router on this particular network.

In the case of equal router priority values, the router ID determines which router becomes the designated router. However, if there already is a designated router on the network when you start this router, it will remain the designated router no matter what your priority or router ID.

By default, each OSPF interface has a router priority of 1.

You can use the BCC or Site Manager to do the following:

- Specify a priority value for the interface.
- Make the router ineligible to be the designated router on this interface.

Using the BCC

To specify a router priority value, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

priority *<priority>*

priority can be 0 (the router is ineligible to become the designated router) or an integer indicating the priority level.

For example, the following command assigns a priority of 2 to interface 2.2.2.2:

```
ospf/2.2.2.2# priority 2
```

Using Site Manager

To specify a router priority value, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Rtr Priority** parameter. Click on **Help** or see the parameter description on page A-44.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Estimating the Transit Delay

By default, OSPF assigns a transmission delay of 1 second to an OSPF interface. You can use the BCC or Site Manager to supply a different transit delay estimate.

Using the BCC

To change the transit delay value, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

transit-delay <*delay*>

delay is the transit delay in seconds.

For example, the following command assigns a transit delay value of 3 seconds to interface 2.2.2.2:

```
ospf/2.2.2.2# transit-delay 3
```

Using Site Manager

To change the transit delay value, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Transit Delay** parameter. Click on **Help** or see the parameter description on page A-44.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the Retransmit Interval

The retransmit interval is the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface.

Each type of network has an optimum retransmit interval. If the interface is connected to a broadcast network, you should use the default setting of 5 seconds. Table 6-3 lists the suggested settings for network types supported by OSPF.

Table 6-3. Retransmit Interval Settings

Network Type	Suggested Retransmit Interval (seconds)
Broadcast	5 (default)
Point-to-point	10
NBMA	10
Point-to-multipoint	10

You can use the BCC or Site Manager to specify a retransmit interval.

Using the BCC

To set the retransmit interval, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

retransmission-interval *<interval>*

interval is the number of seconds between retransmissions.

For example, the following command specifies an OSPF retransmission interval of 10 seconds for IP interface 2.2.2.2:

```
ospf/2.2.2.2# retransmission-interval 10
ospf/2.2.2.2#
```

Using Site Manager

To set the retransmit interval, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Retransmit Interval** parameter. Click on **Help** or see the parameter description on page A-45.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the Hello Interval

The Hello interval specifies how often the router sends Hello messages on the interface. By default, OSPF transmits a Hello message every 10 seconds.

Each type of network has an optimum Hello interval. If the interface is connected to a broadcast network, you should use the default setting -- 10 seconds. Table 6-4 lists the suggested settings for network types supported by OSPF.

Table 6-4. Hello Interval Settings

Network Type	Suggested Hello Interval (seconds)
Broadcast	10 (default)
Point-to-point	15
NBMA	20
Point-to-multipoint	15

Note: This value must be the same for all routers attached to the same network.

You can use the BCC or Site Manager to specify a Hello interval.

Using the BCC

To set the Hello interval, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

hello-interval <*interval*>

interval is the number of seconds between Hello messages.

For example, the following command causes OSPF to transmit a Hello message every 20 seconds on IP interface 2.2.2.2:

```
ospf/2.2.2.2# hello-interval 20
ospf/2.2.2.2#
```

Using Site Manager

To set the Hello interval, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Hello Interval** parameter. Click on **Help** or see the parameter description on page A-45.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the Dead Interval

The dead interval is the number of seconds that OSPF waits to receive a Hello packet from a neighbor before considering the neighbor to be down. The dead interval value should be some multiple of the Hello interval value.

Each type of network has an optimum dead interval. If the interface is connected to a broadcast network, you should use the default setting -- 40 seconds. Table 6-5 lists the suggested settings for network types supported by OSPF.

Table 6-5. Dead Interval Settings

Network Type	Suggested Dead Interval (seconds)
Broadcast	40 (default)
Point-to-point	60
NBMA	80
Point-to-multipoint	60

Note: This value must be the same for all routers attached to the same network.

You can use the BCC or Site Manager to specify a dead interval.

Using the BCC

To set a dead interval, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

dead-interval <*interval*>

interval is the dead interval expressed in seconds.

For example, the following command causes OSPF to wait 60 seconds on IP interface 2.2.2.2 for a Hello message before declaring the neighbor down:

```
ospf/2.2.2.2# dead-interval 60
ospf/2.2.2.2#
```

Using Site Manager

To specify a dead interval, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Dead Interval** parameter. Click on **Help** or see the parameter description on page A-46.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the Poll Interval for NBMA Neighbors

The poll interval is the largest number of seconds allowed between Hello packets sent to an inactive nonbroadcast multi-access (NBMA) neighbor.

By default, each OSPF interface has a poll interval of 120 seconds. You can use the BCC or Site Manager to specify a poll interval.

Using the BCC

To specify a poll interval, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

poll-interval <*interval*>

interval is the poll interval expressed in seconds.

For example, the following command sets the poll interval to 90 seconds on IP interface 2.2.2.2:

```
ospf/2.2.2.2# poll-interval 90
ospf/2.2.2.2#
```

Using Site Manager

To specify a poll interval, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Poll Interval** parameter. Click on **Help** or see the parameter description on page A-46.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the Metric Cost

For OSPF, the best path is the one that offers the least-cost metric.

You must configure cost metrics if you want to specify a preferred path. Allow the preferred path to retain the default cost metric value of 1, and then assign higher-cost metric values to the less-preferred paths.

Figure 6-3 illustrates the benefit of using configurable cost metrics. Assigning the 56 Kb line a cost metric value of 10 forces OSPF to choose the faster T1 line path as the best path, despite the extra hop, when transmitting a packet from host A to host B.

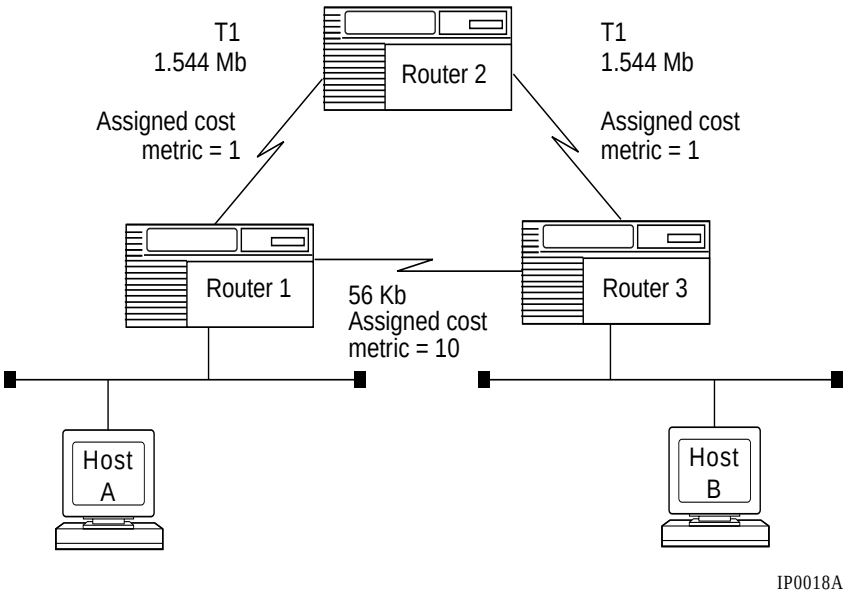

Figure 6-3. Example of Using Configurable Cost Metrics

Table 6-6 lists the suggested values for the metric cost parameter by network type or bit rate.

Table 6-6. Cost Settings

Network Type or Bit Rate	Suggested Metric Cost
100 Mb/s or greater	1 (default)
Ethernet/802.3	10
E1	48
T1	65
64 Kb/s	1562
56 Kb/s	1785
19.2 Kb/s	5208
9.6 Kb/s	10416

By default, each OSPF interface has a cost of 1. You can use the BCC or Site Manager to specify a metric cost for the interface.

Using the BCC

To specify a metric cost, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

metric <*metric*>

metric is the cost of the interface expressed as an integer.

For example, the following command assigns a metric costs of 10 to IP interface 2.2.2.2:

```
ospf/2.2.2.2# metric 10
ospf/2.2.2.2#
```

Using Site Manager

To specify a metric cost, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **Metric Cost** parameter. Click on **Help** or see the parameter description on page A-47.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the MTU Size

OSPF recognizes a maximum transmission unit (MTU) size for updates transmitted on an interface. By default, when you configure OSPF on an interface, OSPF uses the MTU size specified for the type of network to which the interface is connected.

Using the BCC or Site Manager, you can configure OSPF to do the following:

- Send packets no larger than the IP MTU size for Ethernet (1500).
- Use the MTU size that you specify. The number you enter must be less than the IP MTU size for that physical interface.

Note: When running OSPF over a synchronous/PPP link, set the MTU size to a value less than the synchronous MTU size (1200). This setting allows all OSPF routes to be learned over the link.

Using the BCC

To specify the MTU size, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

mtu <*size*>

size is 1, 2, or the MTU size in bytes.

For example, the following command sets the MTU to 2 on IP interface 2.2.2.2. A value of 2 sends packets no larger than the IP MTU size for Ethernet (1500).

```
ospf/2.2.2.2# mtu 2
ospf/2.2.2.2#
```

Using Site Manager

To specify the MTU size, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **MTU Size** parameter. Click on **Help** or see the parameter description on page A-48.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the MTU Mismatch Detect

The mismatch detect controls the interpretation of the MTU field in the database description packet header. According to RFC 2178, the MTU indicates the largest size IP packet that an OSPF interface can receive. If the MTU of a packet is greater than the MTU that the interface can receive, the packet is ignored, and an adjacency is not formed. In RFC 1583, this field does not exist.

Using the BCC

To disable or reenable MTU mismatch detect, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

mtu-mismatch-detect *<state>*

state is one of the following:

enabled (default)
disabled

For example, the following command disables the MTU mismatch detect option on IP interface 2.2.2.2:

```
ospf/2.2.2.2# mtu-mismatch-detect disabled
ospf/2.2.2.2#
```

Using Site Manager

To disable or reenable MTU mismatch detect, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
6. Set the **MTU Mismatch Detect Enable** parameter. Click on **Help** or see the parameter description on page A-49.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a Neighbor on an NBMA Interface

In a nonbroadcast multi-access network, neighbors are not learned dynamically. You can use the BCC or Site Manager to enable and disable the neighbor configuration, supply the IP address for each neighbor, and specify the neighbor's priority. To configure a neighbor, the OSPF interface type must be set to NBMA (see "Specifying the Interface Type" on page 6-21).

Using the BCC

To configure a neighbor on an NBMA interface, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

neighbor *<ip_address>*

ip_address specifies the IP address of the neighbor for this NBMA interface.

To change the default values for a neighbor, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 6-7.

Table 6-7. BCC Neighbor Parameters for an NBMA Interface

Parameter	Values	Function
priority	1 (default) 0 through 255	Specifies a router priority for this NBMA interface. A value of 255 is the highest priority; a value of 0 makes this router ineligible to be the designated router on this network.
state	enabled (default) disabled	Enables or disables this neighbor configuration for this NBMA interface.

Using Site Manager

To configure an OSPF neighbor on an NBMA interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Interfaces**.	The OSPF Interfaces window opens.
5. Click on the OSPF interface to which you want to add a neighbor.	The parameter values for that interface appear in the OSPF Interfaces window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
6. Click on **Neighbors**.	The OSPF Neighbors window opens.
7. Click on **Add**.	The OSPF Neighbor Configuration window opens.
8. Set the **Neighbor's IP Address** parameter. Click on **Help** or see the parameter description on page A-49.	
9. Click on **OK**.	You return to the OSPF Neighbors window.
10. Set the following parameters: • **Enable** • **Priority** Click on **Help** or see the parameter descriptions beginning on page A-50.	
11. Click on **Apply**, and then click on **Done**.	You return to the OSPF Interfaces window.

Deleting OSPF from an IP Interface

To delete OSPF from an interface on which it is currently configured, proceed as follows.

Using the BCC

To delete OSPF from an interface, navigate to an OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

delete

For example, the following command deletes OSPF from interface 2.2.2.2:

```
ospf/2.2.2.2# delete
ip/2.2.2.2/255.255.255.0#
```

Using Site Manager

To delete OSPF from an interface, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on the connector from which you want to delete OSPF services.	The Edit Connector window opens.
2. Click on **Edit Circuit**.	The Circuit Definition window opens.
3. Choose **Protocols**.	The Protocols menu opens.
4. Choose **Add/Delete**.	The Select Protocols window opens. The OSPF button is checked to show that OSPF is enabled on the circuit.
5. Click on **OSPF**.	
6. Click on **OK**.	Site Manager deletes OSPF services from the connector and returns you to the Circuit Definition window.
7. Choose **File**.	The File menu opens.
8. Choose **Exit**.	You return to the Configuration Manager window.

Defining an Area

You define an OSPF area by setting parameters as described under the following topics:

Topic	Page
Supplying an ID for the Area	6-40
Disabling and Reenabling an Area	6-42
Configuring Authentication	6-43
Configuring a Summary Route	6-44
Configuring a Stub Area	6-46
Configuring a Not-So-Stubby Area (NSSA)	6-48

Supplying an ID for the Area

Each OSPF area has a unique identifier. You can use the BCC or Site Manager to supply the ID of the OSPF area that you want to define.

Using the BCC

To create an OSPF area, navigate to an OSPF interface prompt (for example, **box; ip; ospf**) and enter:

area area-id <*area-id*>

area-id is an area identifier in dotted-decimal notation.

To display area parameters and their current values, enter:

info

OSPF displays all area attributes and their current values.

For example, the following command sequence creates OSPF area 0.0.0.0 and displays the default area attributes:

```
ospf# area area-id 0.0.0.0
area/0.0.0.0# info
   area-id 0.0.0.0
   area-type non-stub
   authentication-type none
   stub-metric 1
   import-summaries true
   nssa-translate-to-5 disabled
   nssa-default-originate disabled
   nssa-default-propagate disabled
   nssa-default-ase-path type1
   state enabled
```

Using Site Manager

To create an OSPF area, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Areas**.	The OSPF Areas window opens.
5. Click on **Add**.	The OSPF Area Configuration window opens.
6. Set the **OSPF Area** parameter. Click on **Help** or see the parameter description on page A-51.	
7. Click on **OK**.	You return to the OSPF Areas window.
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Disabling and Reenabling an Area

When you define an OSPF area, the area is automatically enabled. You can use the BCC or Site Manager to disable and reenable the area.

Using the BCC

To disable or reenable an area, navigate to an OSPF area prompt (for example, **box; ip; ospf; area/0.0.0.1**) and enter:

state *<state>*

state is either:

enabled (default)
disabled

For example, this command disables area 0.0.0.1:

```
area/0.0.0.1# state disabled
area/0.0.0.1#
```

Using Site Manager

To disable or reenable an OSPF area, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Areas**.	The OSPF Areas window opens.
5. Click on the area that you want to edit.	The parameter values for that area appear in the OSPF Areas window.
6. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-51.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring Authentication

OSPF provides a measure of security for an area through the use of passwords. If an area is configured to use authentication, all OSPF interfaces configured in that area must be configured with the same password. Different networks can have different passwords.

In such an area, a router that receives a packet verifies the password before it processes the packet. Unauthorized routers are not allowed to communicate with the OSPF system.

By default, authentication is disabled in an area. You can use the BCC or Site Manager to enable authentication and specify a password.

Using the BCC

To enable authentication, navigate to an OSPF area prompt (for example, **box; ip; ospf; area/0.0.0.2**) and enter:

authentication-type simplepassword

After you enable authentication on the area, you can specify a password. Navigate to the OSPF interface prompt (for example, **box; eth 2/2; ip/2.2.2.2/255.255.0.0; ospf**) and enter:

authentication *<string>*

string is any ASCII string up to eight characters long.

Using Site Manager

To enable authentication on an OSPF area and to set a password, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Choose **Areas**.	The OSPF Areas window opens.
5. Click on the area that you want to edit.	The parameter values for that area appear in the OSPF Areas window.
6. Set the **Authentication Type** parameter. Click on **Help** or see the parameter description on page A-51.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.
8. If you set the **Authentication Type** to **Simplepassword**, choose **Protocols**.	The Protocols menu opens.
9. Choose **IP**.	The IP menu opens.
10. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
11. Choose **Interfaces**.	The OSPF Interfaces window opens.
12. Click on the OSPF interface that you want to edit.	The parameter values for that interface appear in the OSPF Interfaces window.
13. Set the **Password** parameter. Click on **Help** or see the parameter description on page A-47.	
14. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a Summary Route

OSPF border routers generate summary advertisements for their attached areas. Each summary advertisement specifies a range of destinations in an area. An area range specification consists of a network address and a variable-length mask. For example, a summary advertisement for the destination 140.191.0.0 with a mask of 255.255.0.0 describes a single route to the collection of destinations 140.191.0.0 through 140.191.255.255. When a packet is forwarded, it is always forwarded to the network that is the best (longest or most specific) match for the packet's destination.

You can use the BCC or Site Manager to configure a summary route.

Using the BCC

To configure a summary route, navigate to an OSPF area prompt (for example, **box; ip; ospf; area/0.0.0.2**) and enter:

summary network *<ip_address>* **mask** *<ip_mask>*

ip_address and *ip_mask* are an IP address/mask pair defining the summary route.

For example, the following command creates the summary route 140.191.0.0 with a mask of 255.0.0.0:

```
area/0.0.0.1# summary network 140.191.0.0 mask 255.0.0.0
summary/0.0.0.1/140.191.0.0#
```

By default, OSPF advertises the summary route. To change the setting, navigate to the summary prompt and enter:

action *<action>*

action is one of the following:

advertise (the default)
block

Using Site Manager

To configure a summary route, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Areas**.	The OSPF Areas window opens.
5. Click on the area for which you want to define a range.	The parameter values for that area appear in the OSPF Areas window.
6. Click on **Ranges**.	The OSPF Ranges window opens.
7. Click on **Add**.	The OSPF Range Area window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
8. Set the following parameters: • **Range Net** • **Range Mask** Click on **Help** or see the parameter descriptions beginning on page A-55.	
9. Click on **OK**.	You return to the OSPF Ranges window.
10. To change the default values for the range, set one or more of the following parameters: • **Enable** • **Mask** • **Status** • **Metric** Click on **Help** or see the parameter descriptions beginning on page A-56.	
11. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a Stub Area

A stub area does not import ASEs and may or may not import internal route summaries. In place of routes to destinations outside the stub, a border router connected to a stub injects a default route advertisement. When an internal router encounters a datagram addressed to a destination outside the stub, the router forwards it to the border router specified in the default route advertisement.

Assume, for example, that the stub area in Figure 1-6 on page 1-16 has been configured to import no internal or external routing information. Border router 8 receives ASEs and internal summaries from its interface to the backbone. However, border router 8 does not forward the ASEs or summaries to the stub. Instead, it injects a default route that internal routers of the stub area use to forward datagrams to destinations beyond the stub.

Using the BCC

To configure an area as a stub, navigate to an OSPF area prompt (for example, **box; ip; ospf; area/0.0.0.2**) and enter:

area-type stub

By default, a border router that injects a default route into a stub area assigns a cost metric of 1 to that default route. To specify a different cost metric, enter:

stub-metric *<cost>*

cost is an integer.

By default, a border router injects network summaries into an attached stub area. To disable this function, enter:

import-summaries false

For example, the following command sequence configures a stub area and assigns a cost metric of 3 to default routes injected into the stub area:

```
ospf/0.0.0.2# area-type stub
ospf/0.0.0.2# stub-metric 3
```

Using Site Manager

To configure an OSPF stub area, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Areas**.	The OSPF Areas window opens.
5. Click on the area that you want to edit.	The parameter values for that area appear in the OSPF Areas window.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
6. Set the following parameters: • **Import AS Extern** • **Stub Default Metric** • **Import Summaries** Click on **Help** or see the parameter descriptions beginning on page A-52.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a Not-So-Stubby Area (NSSA)

An OSPF not-so-stubby area (NSSA) is similar to an OSPF stub area except that, to a limited degree, an NSSA can import AS external routes. Like a stub area, an NSSA consumes less memory and CPU resources by preventing the flooding of AS external link-state advertisements (LSAs) into the area and by using default routing to external destinations. Unlike a stub area, however, an NSSA is more flexible because it can import external routes into an OSPF routing domain.

You must configure type 7 address ranges to allow OSPF NSSA border routers to control the summarization of external routes for the NSSA. For example, the NSSA in Figure 1-6 on page 1-16 has been configured to import type 7 LSAs through NSSA AS boundary router 15. NSSA border router 11 then translates type 7 LSAs with the propagate bit set to type 5 LSAs and floods them into all non-stub areas.

Using the BCC

To configure a not-so-stubby area (NSSA), navigate to the area prompt (for example, **box; ip; ospf; area/0.0.0.1**) and enter:

area-type nssa

To change the default values for NSSA area parameters, enter:

<area_parameter> <value>

area_parameter value is one of the parameter/value pairs listed in Table 6-8.

Table 6-8. BCC Parameters for OSPF NSSA

Parameter	Value	Function
nssa-default-originate	disabled (default) enabled	Enables or disables the origination of a type 7 default route by an NSSA internal AS boundary router
nssa-default-propagate	disabled (default) enabled	Enables or disables the translation of type 7 default route LSAs to type 5 default route LSAs and flooding to all non-stub areas
nssa-default-ase-path	type1 (default) type2	Specifies the path type of the default route for type 7 AS external LSAs
nssa-translate-to-5	disabled (default) enabled	Enables or disables the forced translation of type 7 AS external LSAs to type 5 LSAs by an NSSA area border router (ABR)
stub-metric	1 (default) to 16777215	Assigns a metric value to the type 3 or type 7 default route
import-summaries	true (default) false	Specifies whether summary advertisements are imported to the NSSA

For example, the following command sequence configures area 0.0.0.1 as an NSSA, enables the origination of a type 7 default route by an internal AS boundary router, and enables the translation of type 7 default route LSAs for flooding to all non-stub areas:

```
area/0.0.0.1# area-type nssa
area/0.0.0.1# nssa-default-originate enabled
area/0.0.0.1# nssa-default-propagate enabled
```

Using Site Manager

To configure a not-so-stubby area (NSSA), complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
4. Choose **Areas**.	The OSPF Areas window opens.
5. Click on the area that you want to edit.	The parameter values for that area appear in the OSPF Areas window.
6. Set the following parameters: • **Import AS Extern** • **NSSA Translate** • **NSSA Originate Def Route** • **NSSA Propagate Def Route** • **NSSA Def Route Type** • **Stub Default Metric** • **Import Summaries** Click on **Help** or see the parameter descriptions beginning on page A-52.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring an NSSA Address Range

To implement an NSSA, you must configure at least one type 7 address range for summarization and filtering in the translation of type 7 LSAs to type 5 LSAs. Each type 7 address range consists of an address/mask pair representing multiple external networks. The NSSA border router checks for configured type 7 address ranges and then originates a single type 5 LSA for each range. If a network is not contained in any configured address range, type 7 LSAs are translated to type 5 LSAs directly and flooded to all non-stub areas.

Using the BCC

To configure an address range for an NSSA, navigate to the area prompt (for example, **box; ip; ospf; area/0.0.0.2**) and enter:

nssa-range *<network>/<mask>*

network and *mask* together specify all networks that belong to the NSSA address range. The range mask is not restricted to the natural address class mask.

To change the default values for an OSPF NSSA address range, enter:

<area_parameter> <value>

area_parameter value is one of the parameter/value pairs described in Table 6-9.

Table 6-9. BCC Address Range Parameters for OSPF NSSA

Parameter	Value	Function
action	advertise (default) block	Specifies whether the border router advertises type 5 LSAs for an NSSA address range. An NSSA border router advertises a single external route for each type 7 address range when you set this parameter to advertise.
ase-tag	0 (default) to 16777215	Specifies the value to insert in the external route tag field of translated type 5 LSAs configured for a type 7 address range
state	enabled (default) disabled	Enables or disables this address range

For example, the following command creates an NSSA address range for NSSA 0.0.0.1:

```
area/0.0.0.1# nssa-range 10.0.0.0/16
nssa-range/10.0.0.0/255.255.255.0#
```

Using Site Manager

To configure an address range for an NSSA, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **NSSA Ranges**.	The OSPF NSSA Ranges window opens.
5. Click on **Add**.	The OSPF NSSA Range window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
6. Set the following parameters: • **NSSA Range Net** • **NSSA Range Mask** Click on **Help** or see the parameter descriptions beginning on page A-58.	
7. Click on **OK**.	You return to the OSPF NSSA Ranges window.
8. To change the default values for the range, set one or more of the following parameters: • **Enable** • **Mask** • **Status** • **External Route Tag** Click on **Help** or see the parameter descriptions beginning on page A-59.	
9. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring an Area Border Router

Each OSPF area is connected to the backbone by one or more border routers. A border router can have connections to multiple areas. In Figure 6-4, for example, border router 2 has an interface to a network in area B and an interface to a network in area C.

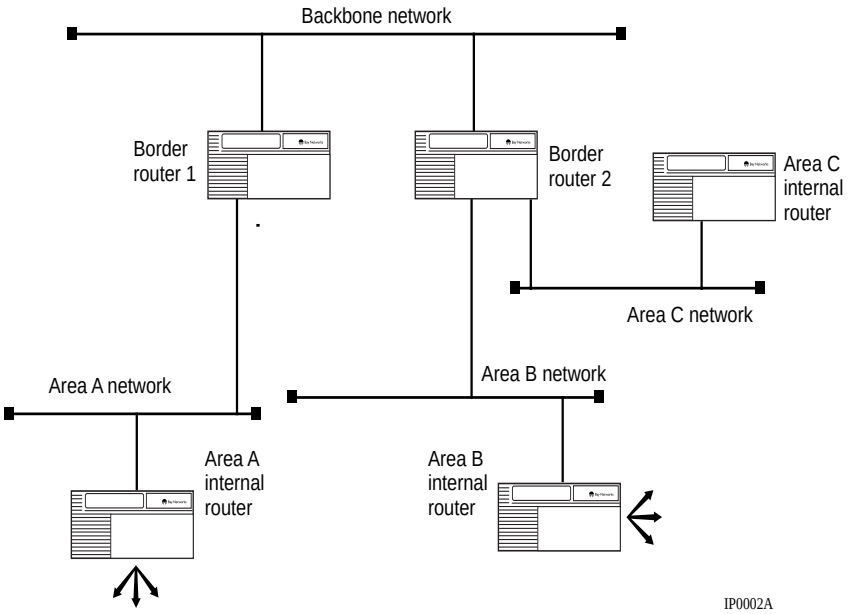

Figure 6-4. Area Border Router

Each border router in an AS does the following:

- Receives routing information from its attached areas, creates summaries of this information, and forwards the summaries to the backbone and to any other attached area. In Figure 6-4, for example, border router 2 floods summaries from area B to the backbone and to area C. Through the backbone, the summaries are forwarded to all other areas in the AS.
- Receives (via the backbone) summaries from other border routers, uses this information to create new routing summaries (which add in the cost of the backbone routes), and forwards the new summaries to its attached areas.

By definition, a border router has an interface to the backbone and interfaces to one or more other areas. To configure OSPF as a border router:

1. **Assign the backbone ID (0.0.0.0) to an OSPF interface.**
2. **Assign an area ID to another OSPF interface.**

Configuring a Virtual Backbone Link Through a Transit Area

Every border router must have a connection to the backbone. This connection can be physical or virtual.

If the border router has an interface to a backbone network, that router is considered to be physically connected to the backbone. In Figure 6-4, border router 1 and border router 2 are both physically connected to the backbone.

In some cases, it may not be possible to configure a border router with an interface to a backbone network. If the router has an OSPF neighbor that is physically connected to the backbone, the router can use that neighbor to establish a *virtual link* to the backbone.

In Figure 6-5, for example, border router 1 has lost its interface to the backbone network. In its place, the network administrator has configured an interface to a network in area B. Through this network, border router 1 now has a neighbor -- border router 2 -- that is connected physically to the backbone. The network administrator can use border router 2 to configure a virtual link between border router 1 and the backbone.

An area that provides a virtual link between a border router and the backbone is called a *transit area*. In Figure 6-5, area B functions as a transit area.

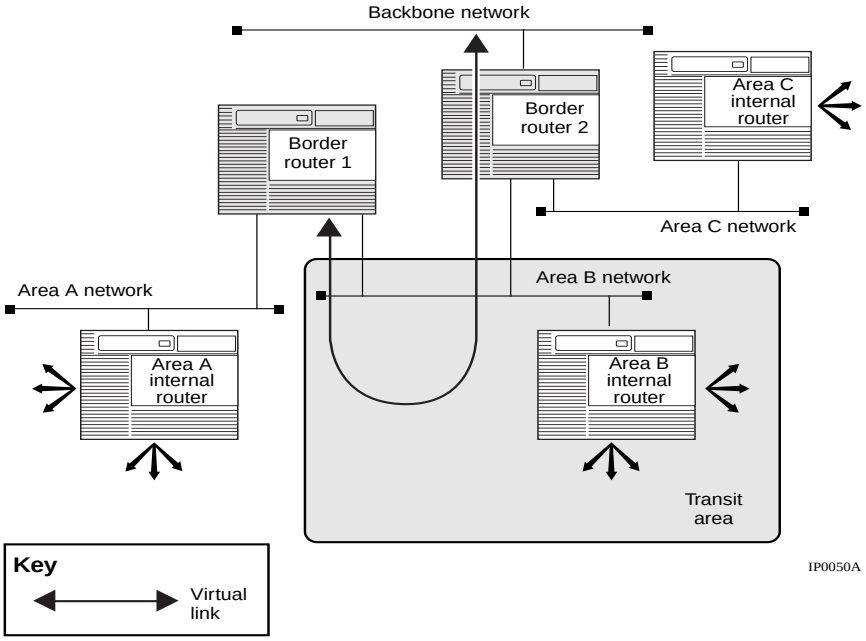

Figure 6-5. Virtual Link and Transit Area

To configure an interface to support a virtual link:

1. **Identify the transit area that supports the virtual link.**
2. **Identify the interface of the OSPF neighbor at the other end of the virtual link.**

After you define the virtual link, you can:

- Enable and disable the virtual link.
- Specify a transit delay, a retransmit interval, a Hello interval, and a dead interval for the link.
- Specify a password.

You can use Site Manager to configure a virtual link.

To configure a virtual backbone link, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF/MOSPF**.	The OSPF/MOSPF menu opens.
4. Choose **Virtual Interfaces**.	The OSPF Virtual Interfaces window opens.
5. Click on **Add**.	The OSPF Virtual Interface Configuration window opens.
6. Set the following parameters: • **Transit Area ID** • **Neighbors Router ID** Click on **Help** or see the parameter descriptions beginning on page A-61.	
7. Click on **OK**.	You return to the OSPF Virtual Interfaces window.
8. To change the default values for the virtual link, set the following parameters: • **Enable** • **Transit Delay** • **Retransmit Interval** • **Hello Interval** • **Dead Interval** • **Password** Click on **Help** or see the parameter descriptions beginning on page A-61.	
9. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring OSPF Accept and Announce Policies

OSPF requires that all routers in a given area maintain a similar routing database. To ensure the integrity of the database, OSPF does not manipulate received link state advertisements before propagating them on an interface.

There are two situations, however, in which an IP policy can be applied to an OSPF interface:

- An OSPF accept policy can be used on a router to control which OSPF non-self-originated external routing information is stored in the routing table. This accept policy controls only what the local router uses; it does not affect the propagation of non-self-originated external information to other routers.
- An OSPF announce policy can be used on a boundary router to control which self-originated external routing updates are placed in the link-state database for distribution according to the OSPF standard. The announce policy affects what other routers learn only with regard to the local boundary router's self-originated information.

When you configure a policy, IP operates according to the following rules:

- IP compares routing information against the *match* criteria in active policies.
- If a match occurs, IP reviews other matching policies for *precedence.*
- IP applies the matching policy with the highest precedence to the routing information and takes the specified *action.*
- IP uses the values of any set criteria in the policy to change the content of the routing information.

The following topics describe how to configure OSPF policies:

Topic	Page
Defining an OSPF Accept Policy	6-58
Specifying Matching Criteria for an OSPF Accept Policy	6-61
Defining an OSPF Announce Policy	6-63
Specifying Modification Values for an OSPF Announce Policy	6-65
Specifying Matching Criteria for an OSPF Announce Policy	6-68

Defining an OSPF Accept Policy

To define a new OSPF accept policy, you must do the following:

- Supply a name for the policy.
- Specify whether OSPF accepts or ignores an update that matches the policy.
- Rank the policy according to preference, precedence, and other criteria.

You can use the BCC or Site Manager to define an OSPF accept policy.

Using the BCC

To define a new OSPF accept policy, navigate to the OSPF global prompt (for example, **box; ip; ospf**) and enter:

accept <*policy_name*>

policy_name is a unique name for the OSPF accept policy.

A policy-specific prompt appears, indicating that the BCC has created the policy using default values for all parameters.

For example, the following command creates an accept policy named pol_1:

```
ospf# accept pol_1
accept/pol_1/ospf#
```

To customize the default values for a policy, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 6-10.

Table 6-10. BCC Definition Parameters for OSPF Accept Policies

Parameter	Values	Function
state	enabled (default) disabled	Enables and disables the policy you created
action	ignore (default) accept	Specifies whether the protocol ignores a route that matches the policy or forwards the route to the routing table manager
preference	1 (default) to 16	Assigns a metric value (the higher the number, the greater the preference) to a route that the protocol forwards to the routing table manager. If confronted with multiple routes to the same destination, the routing table manager may need to use this value to decide which route to insert. Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference, and routes for the most specific networks (longest address and mask) should have the highest preference.
precedence	0 (default) to any integer	Assigns a metric value to this policy (a policy with a higher value takes precedence over a policy with a lower value). This value determines the order of precedence for policies that match the same route.

For example, the following command sets the state to disabled for OSPF accept policy pol_1:

```
accept/pol_1/ospf# state disabled
accept/pol_1/ospf#
```

Using Site Manager

To create an accept policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters.**	The Policy Filters menu opens.
4. Choose **OSPF.**	The OSPF Policies menu opens.
5. Choose **Accept Policies**.	The OSPF Accept Policy Filters window opens.
6. Click on **Add**.	The OSPF Accept IP Policy Filter Configuration window opens.
7. Set the following parameters: • **Name** • **Networks** • **Action** • **Route Preference** • **Rule Precedence** • **Type** • **Tag** Click on **Help** or see the parameter descriptions on pages B-2 through B-4, and pages B-6 and B-7.	
8. Click on **OK**.	You return to the OSPF Accept Policy Filters window.
9. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying Matching Criteria for an OSPF Accept Policy

You can use the BCC or Site Manager to specify matching criteria for the policy.

Using the BCC

To specify matching criteria for an accept policy, navigate to the policy prompt (for example, **box; ip; ospf; accept/pol_1**) and enter:

match

A match prompt appears for the policy.

To specify matching criteria, enter:

<match_criterion> <value>

match_criterion value is one of the pairs listed in Table 6-11.

Table 6-11. BCC Match Parameters for OSPF Accept Policies

Parameter	Values	Function
ase-type	any (default) type1 type2	Describes which type of OSPF ASE route matches this policy
ase-tag*	List of tag values	Specifies OSPF tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain the tag values in this list.
network*	List of IP addresses	Specifies the networks that match this policy. Each specification consists of a network IP address, a mask, and a flag to indicate whether the network address refers to a specific network or to a range of networks. Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Use the default empty list to match any route.

* The BCC implements this matching criterion as an object. To specify multiple elements, create multiple objects.

For example, the following command sequence invokes a match prompt for the OSPF accept policy pol_1 and specifies ASE route type 2:

```
accept/pol_1/ospf# match
match/ospf/accept/pol_1# ase-type type2
```

Using Site Manager

To specify a matching route for an OSPF accept policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters.**	The Policy Filters menu opens.
4. Choose **OSPF.**	The OSPF Policies menu opens.
5. Choose **Accept Policies**.	The OSPF Accept Policy Filters window opens.
6. Select the policy that you want to edit.	
7. Set the following parameters: • **Networks** • **Type** • **Tag** Click on **Help** or see the parameter descriptions on pages B-3 and B-6.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Defining an OSPF Announce Policy

To define a new OSPF announce policy, you must do the following:

- Supply a name for the policy.
- Specify whether OSPF advertises or ignores an update that matches the policy.
- Rank the policy according to precedence and other criteria.

You can use the BCC or Site Manager to define an OSPF announce policy.

Using the BCC

To define a new OSPF announce policy, navigate to the OSPF global prompt (for example, **stack; ip; ospf**) and enter:

announce <*policy_name*>

policy_name is a unique name for the announce policy.

A policy-specific prompt appears, indicating that the BCC has created the policy using default values for all parameters.

For example, the following command creates an OSPF announce policy named pol_1:

```
ospf# announce pol_1
announce/pol_1/ospf#
```

To customize the default values for the policy, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 6-12.

Table 6-12. BCC Definition Parameters for OSPF Announce Policies

Parameter	Values	Function
state	enabled (default) enabled	Enables and disables the policy you created
action	ignore (default) announce	Specifies whether the protocol ignores a route that matches the policy or forwards the route to the routing table manager
precedence	0 (default) to any integer	Assigns a metric value to this policy (a policy with a higher value takes precedence over a policy with a lower value). This value determines the order of precedence for policies that match the same route.

For example, the following command specifies a precedence value of 12 for OSPF announce policy pol_1:

```
announce/pol_1/ospf# precedence 12
announce/pol_1/ospf#
```

Using Site Manager

To create an OSPF announce policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters.**	The Policy Filters menu opens.
4. Choose **OSPF.**	The OSPF Policies menu opens.
5. Choose **Announce Policies**.	The OSPF Announce Policy Filters window opens.
6. Click on **Add**.	The OSPF Announce IP Policy Filter Configuration window opens.

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
7. Set the following parameters: • **Name** • **Action** • **Rule Precedence** Click on **Help** or see the parameter descriptions beginning on page B-8.	
8. Click on **OK**.	You return to the OSPF Announce Policy Filters window.
9. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying Modification Values for an OSPF Announce Policy

You can use the BCC or Site Manager to supply values that OSPF uses to modify a field in an OSPF update that matches the policy.

Using the BCC

To supply a modification value to an update that matches an announce policy, navigate to the policy prompt (for example, **stack; ip; ospf; announce/pol_1**) and enter:

modify

A modification prompt appears for the announce policy.

To specify a modification value, enter:

<parameter> <value>

parameter value is one of the parameter/value pairs listed in Table 6-13.

Table 6-13. BCC Modification Parameters for OSPF Announce Policies

Parameter	Values	Function
ase-tag	0 (default) or a tag value	Specifies a value for the OSPF external route tag field. If the outgoing route matches this policy, OSPF places this value in the field.
ase-type	default type1 type2	Specifies an OSPF ASE metric type to use in advertisements for routes that match this policy
auto-tag	disabled (default) enabled	Enables and disables BGP/OSPF automatic tag generation. Disable auto-tag generation if you want OSPF to use the value that you specify with the ase-tag parameter.
metric	0 (default) or an export metric	Specifies an optional OSPF metric to use when advertising a route that matches this policy
nssa-propagate	disabled (default) enabled	Specifies whether the propagate bit (P-bit) is set in type 7 LSAs of announced OSPF external routes. If the P-bit is set, the NSSA area border router translates type 7 default route LSAs to type 5 default route LSAs before flooding them to all non-stub areas.
advertise*	List of network identifiers	Specifies network IDs to include in place of the network IDs listed in the route to be advertised

* The BCC implements this criterion as an object. To specify a list, create multiple objects.

For example, the following command sequence sets the propagate bit in type 7 LSAs of routes that match OSPF announce policy pol_1:

```
announce/pol_1/ospf# modify
modify/ospf/announce/pol_1# nssa-propagate enabled
```

Using Site Manager

To specify an override value for an OSPF announce policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Select **IP**.	The IP menu opens.
3. Select **Policy Filters.**	The Policy Filters menu opens.
4. Select **OSPF.**	The OSPF Policies menu opens.
5. Select **Announce Policies**.	The OSPF Announce Policy Filters window opens.
6. Select the policy that you want to edit.	
7. Set one or more of the following parameters: • **Advertise** • **OSPF Type** • **OSPF Tag** • **BGP/OSPF Automatic Tag** • **OSPF Metric** • **OSPF NSSA Propagate** Click on **Help** or see the parameter descriptions on page B-11 and on pages B-19 through B-21.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying Matching Criteria for an OSPF Announce Policy

You can use the BCC or Site Manager to specify matching criteria for an OSPF announce policy.

Using the BCC

To specify matching criteria for an OSPF announce policy, navigate to the policy prompt (for example, **stack; ip; ospf; announce/pol_1**) and enter:

match

A match prompt for the policy appears.

For example, the following command invokes a match prompt for the OSPF announce policy pol_1:

```
announce/pol_1/ospf# match
match/ospf/announce/pol_1#
```

To supply matching criteria, enter:

<match_parameter> <value>

match_parameter value is one of the parameter/value pairs listed in Table 6-14.

Table 6-14. BCC Match Parameters for OSPF Announce Policies

Parameter	Values	Function
protocol-source	any (default) direct static rip egp bgp	Specifies one or more route source identifiers. If you select a route source ID, a route from that source that meets the other criteria of this policy matches the policy.
bgp-as*	List of AS numbers	Specifies one or more autonomous system numbers. This policy applies to BGP advertisements received from BGP peers in an AS on this list, and applies only to BGP-sourced routes when BGP is included as a route source.

(continued)

Table 6-14. BCC Match Parameters for OSPF Announce Policies *(continued)*

Parameter	Values	Function
bgp-next-hop*	List of IP addresses	Specifies one or more IP addresses. This policy applies to BGP advertisements whose Next Hop attribute matches an IP address on this list, and applies only to BGP-sourced routes when BGP is included as a route source.
bgp-peer*	List of IP addresses	Specifies the IP address of one or more BGP peers. This policy applies to BGP advertisements authored by a router on this list, and applies only to BGP-sourced routes when BGP is included as a route source.
egp-as*	List of AS numbers	Specifies one or more autonomous system numbers. This policy applies to EGP advertisements received from EGP peers in an AS on this list, and applies only to EGP-sourced routes when EGP is included as a route source.
egp-gateway*	List of IP addresses	Specifies the IP address of one or more EGP gateways. This policy applies to EGP advertisements that use a gateway on this list as the next hop, and applies only to EGP-sourced routes when EGP is included as a route source.
egp-peer*	List of IP addresses	Specifies the IP address of one or more EGP peers. This policy applies to EGP advertisements authored by a router on this list, and applies only to EGP source routes when EGP is included as a route source.
network*	List of IP addresses	Specifies which networks will match this policy. Each identifier consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or to a range of networks. Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Enter an empty list to match any route.

(continued)

Table 6-14. BCC Match Parameters for OSPF Announce Policies *(continued)*

Parameter	Values	Function
rip-gateway*	List of IP addresses	Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list, and applies only to RIP-sourced routes when RIP is included as a route source.
rip-interface*	List of IP addresses	Specifies the address of one or more interfaces on this router. This policy applies to RIP advertisements received on the interfaces in this list, and applies only to RIP-sourced routes when RIP is included as a route source.

* The BCC implements this matching criterion as an object. To specify multiple items, create multiple objects.

Using Site Manager

To specify matching criteria for an OSPF announce policy, complete the following tasks:

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Policy Filters.**	The Policy Filters menu opens.
4. Choose **OSPF.**	The OSPF Policies menu opens.
5. Choose **Announce Policies**.	The OSPF Announce Policy Filters window opens.
6. Select the policy that you want to edit.	

(continued)

Site Manager Procedure *(continued)*	
You do this	**System responds**
7. Set one or more of the following parameters: • **Networks** • **Route Source** • **From RIP Gateway** • **Received on RIP Interface** • **From OSPF Router ID** • **Received OSPF Type** • **Received OSPF Tag** • **From EGP Peer** • **From EGP AS** • **Received EGP Gateway** • **From BGP Peer** • **From BGP AS** • **Received BGP Next Hop** Click on **Help** or see the parameter descriptions beginning on page B-8.	
8. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Appendix A
Site Manager Parameters

This appendix contains the Site Manager parameter descriptions for IP, OSPF, and RIP services. You can display the same information using Site Manager online Help.

This appendix contains the following information:

Topic	Page
IP Parameters	A-2
RARP Parameters	A-35
OSPF Parameters	A-37
RIP Parameters	A-64

IP Parameters

Topic	Page
IP Configuration Parameters	A-2
IP Interface Parameters	A-4
IP Global Parameters	A-19
Static Route Parameters	A-26
Adjacent Host Parameters	A-30
Router Discovery Parameters	A-33

IP Configuration Parameters

Parameter: **IP Address**

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: 0.0.0.0 or any valid IP address

Function: Assigns a 32-bit IP address to the interface.

Instructions: Enter the IP address of the interface in dotted-decimal notation. Enter 0.0.0.0 to configure an unnumbered interface on the circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.4

Parameter: Subnet Mask

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: The Configuration Manager automatically calculates an appropriate subnet mask, depending on the class of the network to which the interface connects. However, you can change the subnet mask with this parameter.

Function: Specifies the network and subnetwork portion of the 32-bit IP address.

Instructions: Either accept the assigned subnet mask or enter another subnet mask in dotted-decimal notation. Enter 0.0.0.0 if you are configuring an unnumbered interface on the circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.6

Parameter: Transmit Bcast Addr

Path: Select IP from the Select Protocols window and click on OK.

Default: 0.0.0.0

Options: 0.0.0.0 or any valid IP broadcast address

Function: Specifies the broadcast address that this IP subnet uses to broadcast packets. Accepting 0.0.0.0 for this parameter specifies that the IP router will use a broadcast address with a host portion of all 1s. Accepting 0.0.0.0 does not configure the router to use the address 0.0.0.0 to broadcast packets. For example, if you have IP address 123.1.1.1 and a subnet mask of 255.255.255.0, accepting the default value 0.0.0.0 configures the IP router to use the address 123.1.1.255 to broadcast packets. To set the explicit broadcast address of all 1s, enter 255.255.255.255 for this parameter.

Instructions: Accept the default, 0.0.0.0, unless the calculated broadcast address (host portion) of all 1s is not adequate. If this is the case, then enter the appropriate IP broadcast address in dotted-decimal notation. If you set the IP Address parameter to 0.0.0.0 (to configure an unnumbered interface), Site Manager automatically sets this parameter to 255.255.255.255.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.8

Parameter: **UnNumbered Assoc Address**

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: Any valid IP address

Function: Specifies an address that IP uses when sourcing a packet. RIP uses this address to make decisions about advertising subnets over the unnumbered interface. RIP advertises subnets over the unnumbered interface if the subnets have the same mask as the associated address.

Instructions: Specify the address of any numbered interface on the router. If you are running RIP over the unnumbered interface and if you are using a subnet address as the associated address, the local and remote associated addresses should have the same network number. If you configure local and remote associated addresses using different network numbers, you must use RIP2 mode.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.110

IP Interface Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Enable

Options: Enable | Disable

Function: Enables or disables IP routing on this interface.

Instructions: Set to Disable to disable IP routing over this circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.2

Parameter: **Mask**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Subnet mask specified when you added IP to the circuit

Options: Depends on the class of the network to which the interface connects

Function: Specifies the network and subnetwork portion of the 32-bit IP address.

Instructions: Enter the subnet mask in dotted-decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.6

Parameter: Broadcast Address

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Broadcast address specified when you added IP to the circuit

Options: 0.0.0.0 or any IP address

Function: Specifies the broadcast address that the IP router uses to broadcast packets. Accepting 0.0.0.0 for the broadcast address specifies that the IP router will use a broadcast address with a host portion of all ones. Accepting 0.0.0.0 does not configure the router to use the address 0.0.0.0 to broadcast packets. For example, if you set the IP address to 123.1.1.1 and the subnet mask to 255.255.255.0, accepting the default value 0.0.0.0 configures the IP router to use the address 123.1.1.255 to broadcast packets. For the explicit broadcast address of all ones, enter 255.255.255.255 for this parameter.

Instructions: Accept the default, 0.0.0.0, unless the calculated broadcast address (host portion) of all ones is not adequate. If this is the case, then enter the appropriate IP broadcast address in dotted-decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.9

Parameter: Cost

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 1

Options: 1 to the value of the RIP diameter (maximum 127)

Function: Sets the cost of this interface. The interface cost is added to routes learned on this interface through RIP and is specified in subsequent RIP packets transmitted out other interfaces.

Instructions: Enter the interface cost value (standard RIP implementation assigns a cost of 1); however, keep in mind that increasing this value causes the upper bound set by the RIP Network Diameter parameter to be attained more rapidly.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.8

Parameter: MTU Discovery

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disabled

Options: Enabled │ Disabled

Function: Specifies whether path MTU (maximum transmission unit) discovery (RFC 1191) is enabled on this interface. When this option is enabled the Do Not Fragment (DF) bit in the IP header is set for all datagrams going out the interface. If any of the datagrams are too large to be forwarded without fragmentation by some router along the path, that router will discard them and return an ICMP destination unreachable message with a code indicating that fragmentation is required and that the DF bit is set. The MTU of the non-conforming hop is also provided. Upon receipt of this message the source host reduces its path MTU for the path. This parameter essentially guarantees that the host will not receive any fragments from a peer doing path MTU discovery.

Instructions: Select Enabled to enable path MTU discovery on this interface; select Disabled to disable the option on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.12

Parameter: ASB

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disabled

Options: Enabled │ Disabled

Function: Specifies whether the IP router floods all-subnet broadcast (ASB) datagrams it receives out this interface. An ASB datagram has a destination address equal to the broadcast address for an entire network (all subnets). For example, if a network interface serves the subnet 128.10.2.1 with a subnet mask of 255.255.255.0, the IP router considers any datagram with a destination address of 128.10.255.255 or 128.10.0.0 to be an ASB datagram.

Instructions: Specify Enabled if you want the IP router to flood ASBs out this interface; specify Disabled to prevent the router from flooding ASBs out this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.14

Parameter: Address Resolution Type

Path: Configuration Manager > Protocols > IP > Interfaces

Default: ARP

Options: ARP │ X.25 DDN │ X.25 PDN │ INARP │ ARPINARP │ NONE │ X.25 BFE DDN │ PROBE │ ARP PROBE │ ATMARP

Function: Specifies the address resolution scheme for this interface. The address resolution scheme enables an originating end station to send broadcast packets containing the network layer address to all nodes on a LAN. The specified end station then sends a unicast packet containing the MAC address to the originating end station.

Instructions: Depending on your network requirements, select one of these options:

- ARP (default)
- X.25_DDN
- X.25_PDN
- INARP only when all frame relay stations support Inverse ARP. INARP enables address resolution for frame relay interfaces, which allows for discovery of the IP address of the remote end of the virtual circuit.
- ARPINARP for your frame relay interfaces. ARPINARP enables both ARP and Inverse ARP.
- NONE to specify no address resolution scheme for this interface
- X.25 BFE DDN
- PROBE to enable HP Probe on the interface
- ARP PROBE to enable both ARP and HP Probe on the interface
- ATMARP to enable ATMARP on the interface (for ATM SVCs configured for classical IP [RFC 1577] and ATM PVCs only)

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.15

Parameter: Proxy

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disabled

Options: Enabled | Disabled

Function: Specifies whether this interface uses Proxy ARP to respond to ARP requests for a remote network.

Instructions: Select Enabled to enable Proxy ARP on this interface. To enable Proxy ARP, you must first set this parameter to Enabled for this interface. When you enable Proxy ARP, the IP router assumes responsibility for IP datagrams destined for the remote network. To enable Proxy ARP for subnets reachable via a default route, also set the global Enable Default Route for Subnets parameter to Enable. To enable Proxy ARP for remote destinations on other networks, set the Nonlocal ARP Destination parameter to Accept. Select Disabled to disable Proxy ARP on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.16

Parameter: Host Cache

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Off

Options: Off | 120 | 180 | 240 | 300 | 600 | 900 | 1200

Function: Specifies whether the IP router times out entries in the address-resolution cache for this interface, and specifies the timeout interval in seconds if the interface does time out entries. The address-resolution cache contains host physical addresses learned by means of ARP or Proxy ARP. A host entry is timed out (deleted) after the specified aging period has elapsed.

Instructions: Select Off to disable timeout on this interface; the IP router does not time out address-resolution cache entries. Select one of the other options to enable timeout with a timeout interval equal to the value you select (for example, 120 seconds); the IP router removes address-resolution cache entries after the specified aging period has elapsed. After an entry is removed, the IP router must use ARP to reacquire the physical layer address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.18

Parameter: Upd Xsum On

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Enabled

Options: Enabled │ Disabled

Function: Specifies whether UDP checksum processing is enabled on this interface.

Instructions: Select Enabled to enable UDP checksum processing for the interface; all outgoing and incoming UDP datagrams are subject to checksumming. You should select Enabled in virtually all instances. Select Disabled to disable UDP checksum processing and provide backward compatibility with UNIX BSD 4.1.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.18

Parameter: MAC Address

Path: Configuration Manager > Protocols > IP > Interfaces

Default: None

Options: 0 │ a user-specified MAC address │ if the interface is on an SMDS circuit, the entire E.164 address -- for example, C1 617 555 5000 FFFF

Function: Specifies a media access control (MAC) address for this IP interface. The IP router will use its IP address and this MAC address when transmitting and receiving packets on this interface.

Instructions: Enter 0 to configure the IP router to use its IP address and the circuit's MAC address when transmitting packets on this interface. Enter your own MAC address to configure the IP router to use its IP address and the specified MAC address when transmitting packets on this interface. If the interface is on an SMDS circuit, by default, IP uses the SMDS-configured address. To configure this parameter for a multinet or multigroup configuration, see *Configuring SMDS*.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.10

Parameter: TR End Station

Path: Configuration Manager > Protocols > IP > Interfaces
Default: Off
Options: On | Off
Function: Specifies source routing over token ring selection.
Instructions: Specify On to enable source route end-station support.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.64

Parameter: TR Endstation ARP Type

Path: Configuration Manager > Protocols > IP > Interfaces
Default: STE
Options: STE | ARE
Function: Specifies the ARP type for an interface configured for token ring support.
Instructions: For spanning tree explorer (STE) ARP packets, use the default. For all routes explorer (ARE) packets, select ARE. Set the TR End Station parameter to On.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.127

Parameter: Interface MTU

Path: Configuration Manager > Protocols > IP > Interfaces
Default: 0
Options: 0 through 65535
Function: Sets the maximum transmission unit (MTU) for an interface to a value less than the MTU of the underlying circuit. If you configure a value greater than the MTU of the underlying circuit, this parameter is ignored. The default value (0) causes the IP interface to use the MTU value of the underlying circuit. This parameter affects the transmission of IP frames only. Transmitted frames that are larger than the MTU value set using this parameter are fragmented into smaller frames. This parameter does not affect frames accepted by the driver.
Instructions: Specify a value (1 through 65535) less than that of the underlying circuit. Otherwise, accept the default value, 0.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.63

Parameter: AMR

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disabled

Options: Enabled │ Disabled

Function: Specifies whether this interface generates Internet Control Message Protocol (ICMP) address-mask reply messages in response to valid address-mask request messages. The interface generates ICMP address-mask reply messages in compliance with the relevant sections of RFCs 950 and 1009.

Instructions: Select Enabled to enable ICMP address-mask reply message generation on this interface. Select Disabled to disable ICMP address-mask reply message generation on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.24.1.13

Parameter: Redirect

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Enable

Options: Enable │ Disable

Function: Indicates whether this interface sends out ICMP redirects. ICMP redirects are messages sent by the router to alert a host that it should be using a different path to route data.

Instructions: Set to Disable if you do not want this interface to send out redirects. For example, in a frame relay network, two stations on the same network may not be directly connected if the network is not fully meshed. Thus, in this case, you would set the Redirect parameter to Disable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.70

Parameter: Ethernet Arp Encaps

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Ethernet

Options: Ethernet │ SNAP │ Both │ Probe (LSAP) │ ARP Enet/Probe LSAP │ ARP SNAP/ Probe LSAP │ ARP Both/Probe LSAP

Function: Specifies the data link encapsulation to use for ARP and HP Probe packets generated at this interface. This parameter is ignored if the underlying media is anything other than Ethernet.

Instructions: Depending on the selection you made for the ARP Resolution Type parameter (ARP, Probe, or ARP/Probe), select the appropriate encapsulation option. If your address-resolution scheme is ARP only, select Ethernet encapsulation, SNAP encapsulation, or Ethernet/SNAP encapsulation. If your resolution scheme is HP Probe only, select LSAP encapsulation. If your resolution scheme is ARP/Probe, select Ethernet/LSAP encapsulation, SNAP/LSAP encapsulation, or Ethernet/SNAP/LSAP encapsulation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.71

Parameter: SMDS Group Address

Path: Configuration Manager > Protocols > IP > Interfaces

Default: SMDS-configured address

Options: A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider

Function: Provides a MAC-layer multicast address for this IP interface in an SMDS network. This parameter is displayed only if this is an SMDS circuit.

Instructions: Enter an entire E.164 address -- for example, E1 617 555 1212 FFFF. If you do not supply an address, IP uses the SDMS-configured address. To configure this parameter for a multinet or multigroup configuration, see *Configuring SMDS.*

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.65

Parameter: SMDS Arp Request Address

Path: Configuration Manager > Protocols > IP > Interfaces

Default: SMDS-configured address

Options: A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider

Function: Provides an address-resolution multicast address for this IP interface in an SMDS network. This parameter is displayed only if this is an SMDS circuit.

Instructions: Enter an entire E.164 address -- for example, E1 617 555 1212 FFFF. If you do not supply an address, IP uses the SMDS-configured address. To configure this parameter for a multinet or multigroup configuration, see *Configuring SMDS*.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.66

Parameter: FRM Broadcast

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 0

Options: Any decimal number

Function: Provides a broadcast address for this IP interface in a frame relay network. If you enter a value for this parameter, the frame relay switch, rather than the router, will broadcast the message. This parameter is displayed only if this is a frame relay circuit.

Instructions: Enter the broadcast address provided in the frame relay subscription agreement.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.67

Parameter: FRM Cast 1 DLCI

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 0

Options: Any decimal number

Function: Provides a multicast address for this IP interface that will send messages to all OSPF routers in a frame relay network. If you enter a value for this parameter, the frame relay switch, rather than the router, will send the message to all OSPF routers. This parameter has meaning only if OSPF has been added to this interface.

Instructions: Enter the multicast address for all OSPF routers as provided in the frame relay subscription agreement.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.68

Parameter: FRM Cast 2 DLCI

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 0

Options: Any decimal number

Function: Provides a multicast address for this IP interface that will send messages to all OSPF designated routers in a frame relay network. If you enter a value for this parameter, the frame relay switch, rather than the router, will send the message to all OSPF designated routers. This parameter has meaning only if OSPF has been added to this interface.

Instructions: Enter the multicast address for all OSPF designated routers as provided in the frame relay subscription agreement.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.69

Parameter:	**Slot Mask**
Path:	Configuration Manager > Protocols > IP > Interfaces
Default:	Slot-mask bit set to 1 (enabling circuitless IP interface support) for every router slot running IP
Options:	For each slot in the router, Site Manager allows you to set the slot-mask bit to 1 (circuitless IP interface support enabled) or 0 (circuitless IP interface support disabled).
Function:	Specifies whether circuitless IP interface support is enabled or disabled on each slot in the router.
Instructions:	If you configured a circuitless IP interface and do not want it to run on certain slots, set the slot-mask bit to 0 on those slots. Be certain to keep the slot-mask bit set to 1 on at least one slot running IP; otherwise, the circuitless IP interface will not initialize. Setting the slot-mask bit parameter to 1 on an empty slot, a slot containing a system resource module, or a slot with no IP support does not affect the circuitless IP interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.4.1.75

Parameter: **Forward Cache Size**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 128

Options: 64-entry minimum; no maximum

Function: Specifies the maximum number of entries allowed in the forwarding table at one time.

Instructions: Specify a forwarding table size for each interface. This parameter controls the number of destinations that are cached in the forwarding table on this receiving interface. When this interface receives an IP packet, the router looks up the destination in the forwarding table. Therefore, an interface that receives packets for a large number of destinations may benefit from a larger forwarding table. The larger the number of entries, the more likely it is that the destination will already be in the forwarding table and the faster the route lookups will be for those destinations. Configuring a forwarding table size that is larger than necessary reduces the total amount of memory usable by other applications. Configuring a forwarding table size that is too small can affect overall router performance. Check the number of cache hits and misses to determine the optimal size of the forwarding table. For debugging purposes, if you see the wfIpInterfaceCacheMisses statistic going up at a rapid rate, consider increasing the table size. However, an occasional cache miss does not warrant an increase in table size.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.104

Parameter: **Enable Security**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disable

Options: Enable | Disable

Function: Specifies whether Revised IP Security Option (RIPSO) is enabled for the interface.

Instructions: If you do not support RIPSO on your network, accept the default setting, Disable. If you are configuring RIPSO support, set this parameter to Enable to access the rest of the RIPSO parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.76

Parameter: **Unnumbered Associated Alternate**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Disable

Options: Enable | Disable

Function: Automatically assigns an alternate associated address to an unnumbered interface in the event that the primary associated address goes down. IP uses the first available interface.

Instructions: Use the alternate unnumbered address option to ensure that the unnumbered interface has a usable associated address on the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.111

Parameter: **ATM ARP Mode**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: Client

Options: Client | Server

Function: Specifies whether the router is running as an ATM client or server on this interface.

Instructions: You must configure one ATMARP server for each logical IP subnet that you define.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.112

Parameter: **ARP Server ATM Address Network Prefix**

Path: Configuration Manager > Protocols > IP > Interfaces

Default: None

Options: *XX*00000000000000000000000 to *XX*FFFFFFFFFFFFFFFFFFFFFFF where *XX* = 39, 45, or 47

Function: Specifies the ATM address network prefix of the ATMARP server on your network.

Instructions: Enter the ATM address network prefix of the ATMARP server on your network. A complete ATM address consists of a network prefix and a user part. Use the ARP Server ATM Address User Part parameter to supply the user part of the ATM address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.113

Parameter: ARP Server ATM Address User Part

Path: Configuration Manager > Protocols > IP > Interfaces

Default: None

Options: 00000000000000 to FEFFFFFFFFFFF

Function: Specifies the user part (suffix) of the ATM address for the ATMARP server on your network. The user part consists of a 6-byte end-station identifier and a 1-byte selector field.

Instructions: Enter the user part suffix of the ATMARP server on your network. A complete ATM address consists of a network prefix and a user part. Use the ARP Server ATM Address Network Prefix parameter to supply the network part of the ATM address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.113

Parameter: Registration Refresh Interval

Path: Configuration Manager > Protocols > IP > Interfaces

Default: 900 seconds for a client; 1200 seconds for a server

Options: Any interval (in seconds)

Function: For a client, this parameter specifies the interval between registration refreshes. For a server, this parameter specifies the duration for which the registration is valid.

Instructions: Determine whether ATMARP is running as a client or as a server on this interface and enter an appropriate value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.115

IP Global Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Global

Default: Enable

Options: Enable │ Disable

Function: Specifies the state of the IP router software.

Instructions: Select Enable if you previously disabled the IP router software and now wish to reenable it. Select Disable to disable the IP router software. In dynamic mode, when you set this parameter to Disable, you immediately prohibit all Site Manager communication with the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.2

Parameter: Forwarding

Path: Configuration Manager > Protocols > IP > Global

Default: Forwarding

Options: Forwarding │ Not Forwarding

Function: Specifies whether the IP router forwards IP traffic that is not explicitly addressed to it.

Instructions: Select Forwarding if you want the IP router to route (forward) IP traffic. Forwarding configures the IP router to process all broadcast packets and all IP packets explicitly addressed to it, and to route all other IP packets. Select Not Forwarding if you want to provide IP management access (by means of TFTP and SNMP) to all active IP interfaces but also want to prohibit the IP router from forwarding IP traffic. You must specify an identical IP address and mask combination for each active IP interface that will provide management access. Not Forwarding configures the IP router to act as an IP host; it does not forward IP traffic, but it still processes packets explicitly addressed to it. In Not Forwarding mode, only static routes and adjacent-host routes are allowed. No routing protocols are initiated. Because the IP router does not forward IP traffic in Not Forwarding mode, you must configure the router to bridge IP traffic not explicitly addressed to it. You must configure the bridge for each circuit that conveys IP datagrams. The bridge then forwards all IP datagrams that are not explicitly addressed to the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.4

Parameter: ARP Forwarding

Path: Configuration Manager > Protocols > IP > Global

Default: Forwarding

Options: Forwarding | Not Forwarding

Function: Specifies how ARP should act in relation to IP's forwarding state. Note that Forwarding means that IP is in forwarding mode. If this parameter is set to Forwarding, then ARP packets are either consumed (if destined for the router) or dropped. If this parameter is set to Not Forwarding, ARP packets are consumed, if destined for the router, or bridged onto remaining ARP interfaces.

Instructions: Always set this parameter the way you set the Forwarding parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.3

Parameter: Nonlocal ARP Source

Path: Configuration Manager > Protocols > IP > Global

Default: Drop

Options: Drop | Drop and Log

Function: Determines what happens when IP encounters an invalid ARP source address. If this parameter is set to Drop and Log, IP logs an invalid ARP source address when processing an ARP request. If this parameter is set to Drop, IP does not log the invalid ARP source address. In either case, IP drops the invalid ARP request.

Instructions: If you want to log the invalid ARP source address, set the parameter to Drop and Log. Otherwise, set the parameter to Drop.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.4

Parameter: Nonlocal ARP Destination

Path: Configuration Manager > Protocols > IP > Global

Default: Drop

Options: Drop | Accept

Function: Determines whether IP drops ARP requests in which the source and destination addresses are located in different networks or subnetworks. This parameter allows Proxy ARP to generate replies when the source and destination networks in the ARP request are different.

Instructions: To process ARP requests with source and destination addresses from different networks, set the parameter to Accept. The Proxy parameter must be set to Enable for the router to generate ARP replies.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.5

Parameter: Default TTL

Path: Configuration Manager > Protocols > IP > Global

Default: 30

Options: 1 to 255 hops

Function: Specifies the starting value of the time-to-live (TTL) counter for each packet that the router originates and transmits (called a source packet). Each router, or hop, that the packet traverses decrements the TTL counter by one. When the counter reaches zero, the router discards the packet unless it is destined for a locally attached network. The TTL counter prevents packets from looping endlessly through the network.

Instructions: Enter the maximum number of hops that a source packet can traverse.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.5

Parameter: RIP Diameter

Path: Configuration Manager > Protocols > IP > Global

Default: 15

Options: 1 to 127

Function: Specifies the value, or hop count, that the Routing Information Protocol (RIP) uses to denote infinity. In order for RIP to operate properly, every router within the network must be configured with an identical RIP diameter value. If RIP is not enabled, this parameter specifies the maximum number of hops within the autonomous system; if RIP is not enabled, the IP router still must understand network width.

Instructions: You must set this parameter so that none of the interface cost, static cost, or route filter cost parameters exceed the RIP diameter. Nortel Networks recommends that you accept the default RIP diameter value of 15.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.6

Parameter: Zero Subnet Enable

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Enable | Disable

Function: Specifies whether an interface address whose subnet portion is all zeros or all ones should be declared legal or not. If you set this parameter to Enable, then you can configure IP interfaces with a subnet ID of zero. Setting this parameter to Disable prevents you from doing so.

Instructions: Accept the default, Disable, if you do not have any interfaces that have a zero subnet ID. Otherwise, reset this parameter to Enable. The use of all-zero subnet addresses is discouraged for the following reason: if an all-zero subnet address and an all-zero broadcast address are both valid, the router cannot distinguish an all-subnets broadcast from a directed broadcast for the zero subnet.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.10

Parameter: Estimated Networks

Path: Configuration Manager > Protocols > IP > Global

Default: 0

Options: 0 to 2,147,483,647

Function: Allows the IP software to preallocate system resources based on the anticipated size of the routing table. Preallocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation. Preallocation also makes better use of memory and reduces the amount of memory required.

Instructions: Set to the number of networks (including unique subnets) that you expect. Avoid using a number that is excessively large. This will cause a wasteful overallocation of memory. If you use the default value, 0, IP software preallocates memory for 500 routing table entries.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.11

Parameter: Estimated Hosts

Path: Configuration Manager > Protocols > IP > Global

Default: 0

Options: 0 to 2,147,483,647

Function: Allows the IP software to preallocate system resources based on the anticipated size of the routing table. Preallocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation.

Instructions: Set to the number of hosts that you expect. Avoid using a number that is excessively large. This will cause a wasteful overallocation of memory. If you use the default value, 0, IP software preallocates memory for 500 routing table entries.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.13

Parameter: **Enable Default Route for Subnets**

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Enable | Disable

Function: Specifies whether the IP router uses a default route for unknown subnets. The default route must be present in the routing table. When you set this parameter to Enable, the IP router uses a default route. When you set this parameter to Disable, the IP router does not use a default route.

Instructions: Accept the default, Disable, if you do not want the IP router to use a default route for unknown subnets. Otherwise, reset this parameter to Enable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.14

Parameter: **Maximum Policy Rules**

Path: Configuration Manager > Protocols > IP > Global

Default: 32

Options: Any integer

Function: Specifies the maximum number of policies that can be configured per policy type (accept or announce) per protocol.

Instructions: To configure more than 32 accept or announce policy rules for a protocol, you must set this parameter to a larger value. IP will round the value up to the next multiple of 32.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.15

Parameter: **Route Filter Support**

Path: Configuration Manager > Protocols > IP > Global

Default: Enable

Options: Enable | Disable

Function: Specifies whether or not IP supports route filters.

Instructions: If you do not require support for route filters, select Disable. Otherwise, accept the default.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.16

Parameter: RIP Maximum Equal Cost Paths

Path: Configuration Manager > Protocols > IP > Global

Default: 1

Options: 1 to 5

Function: Specifies the maximum number of equal-cost paths allowed for a network stored in the routing table by RIP.

Instructions: If you enabled equal-cost multipath support on the router, specify a value from 2 to 5.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.17

Parameter: Multiple Nexthop Calculation Method

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Disable │ Round Robin │ Source Destination Hash │ Destination Hash

Function: Enables and disables equal-cost multipath support for RIP and OSPF and specifies the method that IP uses to choose the next hop when more than one is available. Three methods are available: round-robin selection, selection based on the source addresses (IP forwards all packets with the same source address to the same next hop), and selection based on the source and destination address (IP forwards all packets with the same source and destination address to the same next hop).

Instructions: Configure RIP and OSPF to support equal-cost routes to the same destination.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.18

Parameter: IP OSPF Maximum Path

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: 1

Options: 1 to 5

Function: Specifies the maximum number of equal-cost paths allowed for a network stored in the routing table by OSPF.

Instructions: If you enabled equal-cost multipath support on the router, specify a value from 2 to 5.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.18

Parameter: **Enable ISP Mode Support**

Path: Configuration Manager > Protocols > IP > Global

Default: Disable

Options: Enable | Disable

Function: Enables and disables Internet service provider (ISP) features.

Instructions: Use this parameter to configure BGP as a soloist and to disable the use of forwarding tables on IP interfaces.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.19

Parameter: **Percentage of ARP Buffers**

Path: Configuration Manager > Protocols > IP > Global

Default: 100

Options: An integer indicating the percentage of buffers

Function: Defines the upper limit (as a percentage) of buffers that ARP can use for saving buffers when resolving ARP requests.

Instructions: Specify the percentage as an integer.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.27

Static Route Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > Static Routes

Default: Enable

Options: Enable | Disable

Function: Specifies the state (active or inactive) of the static route record in the IP routing tables.

Instructions: Select Disable to make the static route record inactive in the IP routing table; the IP router will not consider this static route. Select Enable to make the static route record active again in the IP routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.2

Parameter: Destination IP Address

Path:	Configuration Manager > Protocols > IP > Static Routes
Default:	None
Options:	Any valid IP network address
Function:	Specifies the IP address of the network to which you want to configure the static route. Specifies a supernet for which you want to configure a black hole static route.
Instructions:	Enter the destination IP address in dotted-decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet address. You can configure up to 12 static routes to the same destination.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.5.1.3

Parameter: Address Mask

Path:	Configuration Manager > Protocols > IP > Static Routes > **Add**
Default:	None
Options:	Based on the network class of the IP address you specified using the Destination IP Address parameter
Function:	Specifies the subnet mask of the destination network. Specifies the supernet mask of the supernet for which you want to configure a black hole static route.
Instructions:	Enter the subnet or supernet mask in dotted-decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet mask.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.5.1.4

Parameter: Cost

Path: Configuration Manager > Protocols > IP > Static Routes > **Add**

Default: 1

Options: 1 to the value of the RIP Diameter parameter (maximum 127)

Function: Specifies the number of router hops that a datagram can traverse before reaching the destination IP address. The IP router uses the cost value when determining the best route for a datagram to follow.

Instructions: Enter the number of router hops.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.5

Parameter: Next Hop Addr

Path: Configuration Manager > Protocols > IP > Static Routes

Default: 0.0.0.0

Options: Any valid IP address

Function: Specifies the IP address of the next-hop router. Defines a black hole route for a supernet.

Instructions: Enter the IP address in dotted-decimal notation. To configure a black hole static route, enter 255.255.255.255. If you are configuring a static route to an unnumbered interface, enter 0.0.0.0.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.6

Parameter: Next Hop Mask

Path: Configuration Manager > Protocols > IP > Static Routes

Default: 0.0.0.0

Options: Any valid subnet mask address

Function: Specifies the subnet mask of the next-hop router. This parameter also defines a black hole route for a supernet.

Instructions: Enter the subnet mask in dotted-decimal notation. To configure a black hole static route, enter 255.255.255.255. If you are configuring a static route to an unnumbered interface, enter 0.0.0.0.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.7

Parameter: Preference

Path: Configuration Manager > Protocols > IP > Static Routes

Default: 16

Options: 1 to 16

Function: Specifies a weighted value (from 1 to 16, with 16 being the most preferred) that the IP router uses to select a route when its routing tables contain multiple routes to the same destination.

Instructions: Enter a value from 1 to 16 for this static route. To configure a black hole static route, enter the maximum preference value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.8

Parameter: Unnumbered CCT Name

Path: Configuration Manager > Protocols > IP > Static Routes

Default: None

Options: A valid circuit name

Function: Specifies the local router circuit associated with the static route over an unnumbered interface.

Instructions: An entry for a route using an unnumbered interface must include the circuit associated with the interface. Use this parameter to specify that circuit name.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.11

Adjacent Host Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: Enable
Options: Enable | Disable
Function: Specifies the state (active or inactive) of the adjacent host in the IP routing tables.
Instructions: Select Disable to make the adjacent host record inactive in the IP routing table; the IP router will not consider this adjacent host. Select Enable to make the adjacent host record active again in the IP routing table.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.2

Parameter: IP Adjacent Host Address

Path: Configuration Manager > Protocols > IP > Adjacent Hosts > **Add**
Default: None
Options: Any valid IP address
Function: Specifies the IP address of the device for which you want to configure an adjacent host.
Instructions: Enter the IP address in dotted-decimal notation.
MIB Object ID: Not applicable

Parameter: Adjacent Host Address

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: None
Options: Any valid IP address or X.121 address (X.25 only)
Function: Specifies the address of the adjacent host. This address is required for protocols that cannot perform address resolution.
Instructions: Enter the address of the adjacent host. If you configured X.25 as the data link protocol, and the connection type is PDN/X.25, DDN/X.25, or BFE/PVCs, enter the X.121 address of the remote PVC.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.9

Parameter: Next Hop Interface Addr

Path:	Configuration Manager > Protocols > IP > Adjacent Hosts
Default:	0.0.0.0
Options:	Any valid IP address
Function:	Specifies the IP address of the router's network interface to the adjacent host.
Instructions:	Enter the IP address in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.6.1.4

Parameter: MAC/ATM Address, DLCI, VPI/VCI

Path:	Configuration Manager > Protocols > IP > Adjacent Hosts
Default:	None
Options:	Depends on the data link you selected
Function:	Specifies the physical address of the adjacent host. This value can be a 48-bit Ethernet address, a 64-bit SMDS address, an ATM PVC VPI/VCI pair, or, for an ATM SVC, the address of the ATM interface.
Instructions:	Enter the MAC address as a 12-digit hexadecimal number (canonical format), a 32-digit hexadecimal number (SMDS), a VPI/VCI pair *(*for example, 0/32), or a 40-digit ATM address (for example, 39000000000000000000000000000A20037B801).
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.6.1.6

Parameter: Host Encapsulation

Path:	Configuration Manager > Protocols > IP > Adjacent Hosts
Default:	Ethernet
Options:	Ethernet | SNAP | PDN | DDN | SNAPIP | NULL
Function:	Specifies the adjacent host's encapsulation method.
Instructions:	Select Ethernet or SNAP if you are defining a point-to-point network interface or if the adjacent host resides on an Ethernet. For an X.25 interface, select PDN or DDN. For an adjacent host on an ATM logical IP subnet, select SNAP. (SNAPIP and NULL also specify host encapsulation methods for ATM networks.)
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.6.1.7

Parameter: Remote Party Sub-Address

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: None
Options: An SVC subaddress
Function: Specifies the subaddress used to establish an SVC to the adjacent host.
Instructions: Supply a valid SVC subaddress.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.10

Parameter: Remote Party Type of Number

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: International
Options: International | Unknown
Function: Specifies the type of number used to establish an SVC to the adjacent host.
Instructions: Supply the required value.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.11

Parameter: Adjacent Host Type

Path: Configuration Manager > Protocols > IP > Adjacent Hosts
Default: Default
Options: FRE 164 | Default | FRX 121 | FRDLCI
Function: Specifies the type of adjacent host.
Instructions: Supply a value.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.12

Router Discovery Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: Enable
Options: Enable | Disable
Function: Disables and enables Router Discovery on this interface.
Instructions: If you configured this interface with Router Discovery, use this parameter to disable Router Discovery.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.2

Parameter: Broadcast Type

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: Multicast
Options: Multicast | Local | Direct
Function: Specifies the type of broadcast to use in sending advertisements.
Instructions: Use Multicast wherever possible, that is, on any link where all listening hosts support IP multicast.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.5

Parameter: Minimum Interval

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 450
Options: 3 to the value of the Maximum Interval parameter
Function: Specifies the minimum time interval in seconds between advertisements.
Instructions: Specify a value that is no less than 3 seconds and less than the value you set for the Maximum Interval parameter.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.6

Parameter: Maximum Interval

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 600
Options: 4 to 1800
Function: Specifies the maximum time interval in seconds between advertisements.
Instructions: Specify a value from 4 through 1800 seconds that is greater than the value you specified for the Minimum Interval parameter, and is no greater than 1800 seconds.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.7

Parameter: Lifetime

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 1800
Options: A value specifying the number of seconds
Function: Specifies the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts, in the absence of further advertisements.
Instructions: Specify a value that is no less than the value you set for the Maximum Interval parameter and no greater than 9000 seconds.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.8

Parameter: Interface Preference

Path: Configuration Manager > Protocols > IP > Router Discovery
Default: 0
Options: A numeric value
Function: Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet.
Instructions: Enter a value indicating the relative preference of the router address. Enter a preference value of 0x80000000 to indicate to neighboring hosts that the address is not to be used as a default route.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.9

RARP Parameters

Parameter:	**Enable/Disable**
Path:	Configuration Manager > Protocols > IP > Reverse ARP > Globals
Default:	Enable
Options:	Enable | Disable
Function:	Globally enables or disables RARP on all interfaces.
Instructions:	Set to Disable if you want to disable RARP for the entire router. Set to Enable if you previously disabled RARP on the router and now want to reenable it.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.9.1.2

Parameter:	**Enable/Disable**
Path:	Configuration Manager > Protocols > IP > Reverse ARP > Interface Table
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables RARP on this interface.
Instructions:	Select Enable to reenable a disabled interface. Select Disable to disable an interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.9.3.1.2

Parameter:	**MAC Address**
Path:	Configuration Manager > Protocols > IP > Reverse ARP > Map Table > **Add**
Default:	None
Options:	Any valid MAC address
Function:	Specifies the MAC address of a client that will use the RARP services of this router.
Instructions:	Enter the MAC address of a client.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.9.2.1.2

Parameter:	**IP Address**
Path:	Configuration Manager > Protocols > IP > Reverse ARP > Map Table > **Add**
Default:	0.0.0.0
Options:	Any valid IP address
Function:	Specifies the corresponding IP address for the client with the MAC address that you specified using the MAC Address parameter. When the router receives a RARP request from the client, the router assigns this IP address to the client and includes it in a response packet.
Instructions:	Enter the IP address corresponding to the value of the MAC Address parameter. Do not accept the default value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.9.2.1.3

OSPF Parameters

Topic	Page
OSPF Global Parameters	A-37
OSPF Interface Parameters	A-42
Neighbor Parameters for an NBMA Interface	A-49
OSPF Area Parameters	A-51
Area Range Parameters	A-55
OSPF NSSA Range Parameters	A-58
OSPF Virtual Interface Parameters	A-61

OSPF Global Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: Enable

Options: Enable | Disable

Function: Globally enables or disables OSPF on all router interfaces.

Instructions: Set to Disable if you want to disable OSPF for the entire router. Set to Enable if you previously disabled OSPF on the router and now want to reenable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.2

Parameter: Router ID

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: The IP address of the first OSPF circuit configured on this router

Options: Any IP address; preferably, an IP interface address on the router.

Function: This IP address uniquely identifies this router in the OSPF domain. By convention, and to ensure uniqueness, one of the interface addresses for the router should be used as the router ID. The router ID will determine the designated router on a broadcast link if the priority values of the routers being considered are equal. The higher the router ID, the greater its priority.

Instructions: Enter the appropriate IP address in dotted-decimal notation. If both OSPF and BGP are running on the router, the OSPF router ID must be identical to the BGP identifier. In addition, the OSPF router ID must match one of the IP addresses configured on the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.4

Parameter: AS Boundary Router

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: No

Options: Yes | No

Function: Indicates whether or not this router functions as an AS boundary router. Only AS boundary routers are allowed to convert non-OSPF routes into OSPF routes so that they can be passed along throughout the OSPF routing domain. The router can be an AS boundary router if one or more of its interfaces is connected to a non-OSPF network (for example, RIP, BGP, or EGP).

Instructions: Set this parameter to Yes if this router functions as an AS boundary router. Otherwise, accept the default value, No.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.7

Parameter: Hold Down Timer

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Global
Default:	1 second
Options:	0 to 10 seconds
Function:	Prevents the algorithm that computes a route to not run more than once per holddown time. Its purpose is to free up the CPU. A value of 0 means there is no holddown time.
Instructions:	Either accept the default value of 1 second or enter a new value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.9

Parameter: OSPF Slot

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Global
Default:	All slots
Options:	Any slot on the router
Function:	Indicates which slots the OSPF soloist is eligible to run on. If the slot on which the OSPF soloist is running goes down, the router will attempt to run OSPF on another slot specified by this parameter.
Instructions:	Select all of the appropriate slots. Use caution when selecting the slots on which OSPF may run. If you choose an empty slot, and it is the only slot you choose, OSPF will not run; if you choose a slot that becomes disabled, and it is the only slot you choose, OSPF will not restart.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.10

Parameter: ASE Metric Support

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Global
Default:	Disable
Options:	Enable | Disable
Function:	Causes the router to use the route weight as the OSPF metric in OSPF ASE type 2 advertisements.
Instructions:	Disable ASE metric support if the router is to interoperate with routers using an OSPF version earlier than Version 8.00. The new metric is not compatible with the earlier metric.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.11

Parameter: Backup Enable

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Global
Default:	Disable
Options:	Enable | Disable
Function:	Enables or disables the backup link state database of the OSPF backup soloist. When the parameter is set to Disable, the OSPF backup soloist will not maintain a copy of the OSPF link state database.
Instructions:	Select the default, Disable, if you do not want to back up the OSPF soloist.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.12

Parameter: Primary Log Mask

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Global
Default:	Trace | Info | Debug | INTF state | NBR state | Bad LS
Options:	Trace | Info | Debug | INTF state | NBR state | LSA self-origin | LSA receipt | Route change | Bad LS | Less recent LSA | More recent LSA | Max age LSA
Function:	Specifies which OSPF log messages should be logged in the primary log.
Instructions:	Highlight the value of the Primary Log Mask parameter and click on Values. Choose the log messages that you want to enter into the primary log, then click on OK.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.13

Parameter: Backup Log Mask

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Global
Default:	Log no messages
Options:	Trace | Info | Debug | INTF state | NBR state | LSA self-origin | LSA receipt | Route change | Bad LS | Less recent LSA | More recent LSA | Max age LSA
Function:	Specifies which OSPF log messages should be logged in the backup log.
Instructions:	Highlight the value of the Backup Log Mask parameter and click on Values. Choose the log messages that you want to enter into the backup log, then click on OK.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.1.14

Parameter: **Tag Generation Method**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: Zero

Options: Zero | Autotag | Proprietary

Function: Specifies the method of OSPF external tag field generation.

Instructions: Set the parameter to Autotag if you want OSPF to generate a tag value according to RFC 1403, *OSPF/BGP Interaction.* Use the default to insert 0 into the tag field. The Proprietary option is reserved for debugging purposes.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.15

Parameter: **Multicast Extensions**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: 0

Options: 0 (no multicast forwarding is enabled) | 1 (intra-area multicasting only) | 3 (intra-area and inter-area multicasting) | 5 (intra-area and inter-AS multicasting) | 7 (multicasting everywhere)

Function: Indicates whether the router is forwarding IP multicast (Class D) datagrams based on the algorithms defined in the Multicast Extensions to OSPF.

Instructions: Set the bitmask as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.21

Parameter: **Timeout Value**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: 600 seconds

Options: An integer

Function: Specifies a timer value for timing out MOSPF forwarding entries.

Instructions: Use the default setting.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.26

Parameter: **RFC 1583 Compatibility Enable**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Global

Default: Enabled

Options: Enabled │ Disabled

Function: Controls the preference rules used when choosing among multiple AS external LSAs advertising the same destination.

Instructions: Set this parameter to Enabled to use the preference rules specified by RFC 1583. Set this parameter to Disabled to use the preference rules specified in RFC 2178, which prevent routing loops when AS external LSAs for the same destination originate from different areas.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.29

OSPF Interface Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces

Default: Enable

Options: Enable │ Disable

Function: Indicates whether or not OSPF is enabled on this interface. The default value, Enable, indicates that neighbor relationships may be formed on this interface, and that this interface will be advertised as an internal route to some area. The value Disable indicates that this is not an OSPF interface.

Instructions: Set this parameter to Disable if you do not want OSPF enabled on the interface. Set it to Enable if you previously disabled OSPF on this interface and now want to reenable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.2

Parameter: Area ID

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	0.0.0.0
Options:	Any 4-octet number in dotted-decimal notation
Function:	Identifies the area to which this interface belongs.
Instructions:	Enter the appropriate area ID in dotted-decimal notation. Area ID 0.0.0.0 is reserved for the backbone.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.6

Parameter: Type

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	Broadcast
Options:	Broadcast │ NBMA (nonbroadcast multiaccess) │ Point-to-point │ Point-to-multipoint │ Point-to-multipoint (STD) │ Passive
Function:	Indicates the type of network to which this interface is attached. Choose Broadcast if this network is a broadcast LAN, such as Ethernet. Choose NBMA if the network is a nonbroadcast network, such as X.25. Choose Point-to-point for a synchronous, point-to-point interface. If you want to use the Nortel Networks proprietary solution for frame relay point-to-multipoint networks, choose Point-to-multipoint. Choose Point-to-multipoint (STD) if the network is a point-to-multipoint network. Choose Passive to configure an interface that OSPF cannot use to form neighbor relationships. OSPF cannot accept Hello messages or send advertisements on the passive interface.
Instructions:	Set this parameter to match this interface type. If you set this parameter to NBMA, you must configure neighbors manually.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.7

Parameter: Rtr Priority

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces

Default: 1

Options: 0 to 255

Function: Indicates the priority of this interface. The router priority value is used in multiaccess networks (broadcast, NBMA, or point-to-multipoint), for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network. In the case of equal Rtr Priority values, the router ID determines which router becomes the designated router. However, if there already is a designated router on the network when you start this router, it will remain the designated router no matter what your priority or router ID.

Instructions: Set the router priority to a value from 0 through 255 or accept the default value, 1.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.8

Parameter: Transit Delay

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces

Default: 1 second

Options: 1 to 3600 seconds

Function: Indicates the estimated number of seconds it takes to route a packet over this interface.

Instructions: Either accept the default value of 1 second or enter a slightly higher number for slower-speed serial lines, for example, 15 to 20 seconds for a 19.8 KB line.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.9

Parameter: Retransmit Interval

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	5 seconds
Options:	1 to 3600 seconds
Function:	Indicates the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. This value is also used when retransmitting OSPF packets. Nortel Networks suggests the following values for this parameter: for broadcast, 5 seconds; for point-to-point, NBMA, and point-to-multipoint, 10 seconds.
Instructions:	Either accept the default value of 5 seconds or set the retransmit interval to a slightly higher number for slower-speed serial lines.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.10

Parameter: Hello Interval

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	10 seconds
Options:	1 to 65,535 seconds
Function:	Indicates the number of seconds between the Hello packets that the router sends on the interface. Nortel Networks suggests the following values for this parameter: for broadcast, 10 seconds; for point-to-point, 15 seconds; for NBMA, 20 seconds; for point-to-multipoint, 15 seconds.
Instructions:	Either accept the default value of 10 seconds or set the Hello interval to a higher number for slower-speed serial lines. This value must be the same for all routers attached to the same network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.11

Parameter:	**Dead Interval**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	40 seconds
Options:	1 to 2,147,483,647 seconds
Function:	Indicates the number of seconds that a router's Hello packets have not been seen before its neighbors declare the router down. The dead interval value should be some multiple of the Hello interval value. Nortel Networks suggests the following values for this parameter: for broadcast, 40 seconds; for point-to-point, 60 seconds; for NBMA, 80 seconds; for point-to-multipoint, 60 seconds.
Instructions:	Either accept the default value of 40 seconds or set the dead interval to a higher number for slower-speed serial lines. This value must be the same for all routers attached to the same network
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.12

Parameter:	**Poll Interval**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	120 seconds
Options:	1 to 2,147,483,647 seconds
Function:	Indicates the largest number of seconds allowed between Hello packets sent to an inactive NBMA neighbor.
Instructions:	Either accept the default value of 120 seconds or set this parameter to a slightly higher number for slower-speed serial lines.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.13

Parameter: Metric Cost

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	1
Options:	1 to 65,535
Function:	Indicates the cost of using this type of service on this interface. Nortel Networks suggests the following values for this parameter: for 100 Mb/s or greater, 1; for Ethernet/802.3, 10; for E1, 48; for T1, 65; for 64 Kb/s, 1562; for 56 Kb/s, 1785; for 19.2 Kb/s, 5208; for 9.6 Kb/s, 10416. This parameter allows you to configure preferred paths. If you do want to configure a preferred path, allow that path to retain the default value of 1 or assign it a relatively low metric cost. Then, assign the less preferred paths a higher metric cost value.
Instructions:	Either accept the default value, 1, or enter a larger number for a slower path or a backup route.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.16

Parameter: Password

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	None
Options:	Any ASCII string up to eight characters long
Function:	Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. This parameter is valid only when the Authentication Type parameter for the area is set to Simplepassword.
Instructions:	Enter the appropriate password. All routers in the same area must either have no authentication or have the same password.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.17

Parameter: MTU Size

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces

Default: 1

Options: 1 | 2 | a number up to 10,000

Function: Specifies the maximum transmission unit (MTU) size of OSPF updates on this interface.

Instructions: Accept the default value, 1, to use the IP MTU size for that physical interface. Enter 2 to send packets no larger than the IP MTU size for Ethernet (1500). Enter a number up to 10,000 to specify an MTU size directly; the number you enter must be less than the IP MTU size for that physical interface. When running OSPF over a synchronous/PPP link, set the MTU size to a value less than the synchronous MTU size (1200). This setting allows all OSPF routes to be learned over the link.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.29

Parameter: Multicast Forwarding

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces

Default: Blocked

Options: Blocked | Multicast | Unicast

Function: Specifies the way multicasts should be forwarded on this interface: not forwarded, forwarded as data link multicasting, or forwarded as data link unicasts. Data link multicasting is not meaningful on point-to-point and NBMA interfaces.

Instructions: If you configured MOSPF globally, specify the way that you want IP to forward multicast packets on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.30

Parameter:	**MTU Mismatch Detect Enable**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces
Default:	Enabled
Options:	Enabled | Disabled
Function:	Controls the interpretation of the MTU field in the database description packet header. According to RFC 2178, the MTU indicates the largest size IP packet that an OSPF interface can receive. If the MTU of a packet is greater than the MTU that the interface can receive, the packet is ignored, and an adjacency is not formed. In RFC 1583 this field does not exist.
Instructions:	The value Enabled specifies RFC 2178 processing of the MTU field. The value Disabled specifies RFC 1583 processing.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.5.1.34

Neighbor Parameters for an NBMA Interface

Parameter:	**Neighbor's IP Address**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces > Neighbors > **Add**
Default:	None
Options:	IP address of the neighbor
Function:	Specifies the IP address of an NBMA neighbor for this interface.
Instructions:	Enter the appropriate IP address in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.7.1.4

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces > Neighbors

Default: Enable

Options: Enable | Disable

Function: Allows you to enable and disable this neighbor configuration for this interface. This parameter is useful if you want to temporarily disable a neighbor configuration rather than delete it.

Instructions: Set to Disable if you want to disable this neighbor configuration. Or set to Enable if you previously disabled this neighbor configuration and now want to reenable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.2

Parameter: Priority

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Interfaces > Neighbors

Default: 1

Options: 0 to 255

Function: Indicates the priority of this neighbor, with 255 indicating the highest priority. The neighbor priority value is used in multiaccess networks for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network.

Instructions: Either accept the default neighbor priority value or enter another value from 0 through 255.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.9

OSPF Area Parameters

Parameter: **OSPF Area**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas > **Add**

Default: 0.0.0.0

Options: Any four-octet number in dotted-decimal notation

Function: Identifies the OSPF area.

Instructions: Enter the appropriate area ID in dotted-decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.1

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: Enable

Options: Enable | Disable

Function: Allows you to enable and disable this area. This parameter is useful if you want to temporarily disable an area rather than delete it.

Instructions: Set this parameter to Disable if you want to disable this area. Set this parameter to Enable if you previously disabled the area and now want to reenable it. This action will cause OSPF to restart.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.2

Parameter: **Authentication Type**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: None

Options: None | Simplepassword

Function: Enables or disables password authentication for the area. If you select Simplepassword (enabling password authentication), only those routers that share the correct password will be able to communicate with each other. If you accept the default, None, password authentication is disabled for this area.

Instructions: Either accept the default value, None, to disable password authentication or select Simplepassword to enable password authentication.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.5

Parameter: Import AS Extern

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: Import External

Options: Import External │ Import No External │ Import NSSA

Function: Indicates whether or not an area imports AS external LSAs. If an area imports type 5 AS external LSAs, it is a normal (non-stub) area. If an area does not import AS external LSAs, it is a stub area. If an area imports type 7 external LSAs only, it is an NSSA (not-so-stubby area).

Instructions: Set to Import No External if you want to configure a stub area. Set to Import NSSA if you want to configure a not-so-stubby area (NSSA). Otherwise, accept the default, Import External, to configure a normal area.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.2.1.6

Parameter: Stub Default Metric

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: 1

Options: 1 to 16777215

Function: When an area border router (ABR) is connected to a stub area, it generates a type 3 summary LSA to the area specifying the default route. When an ABR is connected to an NSSA, it generates a type 3 summary LSA if type 3 LSAs are not imported to the NSSA. However, if type 3 LSAs are imported to the NSSA, the ABR generates type 7 LSAs specifying the default route. The stub metric is the cost of the type 3 or type 7 default route. This parameter has meaning only when the Import AS Extern parameter is set to Import No External or Import NSSA.

Instructions: Either accept the stub metric default value, 1, or supply the appropriate stub metric value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.2.1.7

Parameter: Import Summaries

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: True

Options: True | False

Function: Specifies whether summary advertisements are imported by an NSSA or stub area. This parameter should not be set to True unless the Import AS Extern parameter is set to Import No External or Import NSSA.

Instructions: Set to False if you do not want to import summary advertisements to the area. Otherwise, accept the default, True.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.2.1.8

Parameter: Cost for PtP Links

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: Enable

Options: Enable | Disable

Function: Indicates whether OSPF calculates the cost for a point-to-point link.

Instructions: Enable or disable cost calculation as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.10

Parameter: NSSA Translate

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: Disable

Options: Enable | Disable

Function: Enables or disables the forced translation of type 7 AS external LSAs to type 5 LSAs by an NSSA area border router (ABR). If you set this parameter to Enable, the ABR translates the LSAs. If you set this parameter to Disable, the ABR becomes a candidate for NSSA translator election if no other ABRs are configured as unconditional NSSA translators. The router with the highest OSPF router ID is designated as the translating router.

Instructions: Set to Enable if you want to force the translation of type 7 AS external LSAs to type 5 LSAs by an NSSA ABR. Otherwise, accept the default, Disable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.2.1.11

Parameter: NSSA Originate Def Route

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: Disable

Options: Enable | Disable

Function: Enables or disables the origination of a type 7 default route by an NSSA internal AS boundary router.

Instructions: Set to Enable if you want the AS boundary router to originate a type 7 default route. Otherwise, accept the default, Disable. You must set this parameter to Enable to enter values for the NSSA Propagate Def Route and NSSA Def Route Type parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.2.1.13

Parameter: NSSA Propagate Def Route

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas

Default: Disable

Options: Enable | Disable

Function: Specifies whether the propagate bit is set in the type 7 default route LSAs originated by an internal AS boundary router. If this parameter is set to Enable, type 7 default route LSAs can be translated into type 5 default route LSAs and flooded to all non-stub areas. If the propagate bit is not set in a type 7 LSA, the translator ignores the LSA.

Instructions: Set to Enable if you want the propagate bit set in type 7 default route LSAs. Otherwise, accept the default, Disable. To use this parameter, you must first set the NSSA Originate Def Route parameter to Enable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.2.1.14

Parameter:	**NSSA Def Route Type**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas
Default:	Type 1
Options:	Type 1 | Type 2
Function:	Specifies the path type of the default route for type 7 AS external LSAs.
Instructions:	Select Type 2 if you want to set the path type of the default route to type 2. Otherwise, accept the Type 1 default route type. To use this parameter, you must first set the NSSA Originate Def Route parameter to Enable.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.2.1.15

Area Range Parameters

Parameter:	**Range Net**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas > **Ranges** > **Add**
Default:	None
Options:	Any network IP address
Function:	Allows you to assign a single network address to a group of subnets. This network address, together with the subnet mask you provide, specifies the subnets to be grouped in this area range. Only one link summary advertisement will be generated for all subnets in this range, rather than one link summary advertisement for each of the subnets included in that network.
Instructions:	Enter the appropriate network IP address in dotted-decimal notation.
MIB Object ID:	1.3.6.14.1.18.3.5.3.2.3.4.1.5

Parameter: **Range Mask**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas > **Ranges** > **Add**

Default: None

Options: Any subnet mask

Function: This parameter, together with the Range Net parameter, specifies all of the networks that belong to this range. The range mask is not restricted to the natural address class mask for the address supplied in the Range Net parameter.

Instructions: Enter the appropriate subnet mask in dotted-decimal notation.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.6

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas > **Ranges**

Default: Enable

Options: Enable │ Disable

Function: Enables or disables this range for the specified area. This parameter is useful if you want to disable the range, rather than delete it.

Instructions: Set this parameter to Disable if you want to disable this range. Set the parameter to Enable if you previously disabled this range and now want to reenable it.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.2

Parameter: **Mask**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas > **Ranges**

Default: None

Options: Any address mask

Function: This parameter allows you to change the mask portion of this area range. Mask, together with Range Net, indicates all of the networks that belong to this range. Mask is not restricted to the natural address class mask for the address supplied in the Range Net parameter.

Instructions: Enter the appropriate address mask in dotted-decimal notation.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.6

Parameter: Status

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas > **Ranges**
Default: Advertise
Options: Advertise | Do Not Advertise
Function: Specifies whether the border router advertises a summary route to other areas.
Instructions: Select Do Not Advertise if you want to hide the existence of certain networks from other areas. By default, the border router advertises a single route for the range of routes you specify.
MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.7

Parameter: Metric

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Areas > **Ranges**
Default: 0
Options: 0 to 2,147,483,647
Function: Specifies the metric to advertise into other areas as the distance from the OSPF router to any network in the range.
Instructions: If you select 0, the router uses the value calculated by OSPF.
MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.8

OSPF NSSA Range Parameters

Parameter: **NSSA Range Net**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > NSSA Ranges > **Add**

Default: None

Options: Any network IP address

Function: Allows you to assign a single network address to a group of NSSA type 7 subnets. This network address, together with the NSSA range mask, specifies the type 7 subnets to be grouped in this NSSA range. Only one type 5 LSA will be generated for all type 7 subnets in this range, rather than one type 5 LSA for each type 7 subnet imported by the NSSA.

Instructions: Enter the appropriate network IP address in dotted-decimal notation.

MIB Object ID: Not applicable

Parameter: **NSSA Range Mask**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > NSSA Ranges > **Add**

Default: None

Options: Any subnet mask

Function: This parameter, together with the NSSA Range Net parameter, specifies all networks that belong to this range. The range mask is not restricted to the natural address class mask for the address supplied in the NSSA Range Net parameter.

Instructions: Enter the appropriate subnet mask in dotted-decimal notation.

MIB Object ID: Not applicable

Parameter:	**Enable**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > NSSA Ranges
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables this address range for type 7 AS external LSAs within the NSSA. You can use this parameter to temporarily disable an NSSA range.
Instructions:	Set to Disable if you want to disable this NSSA address range. Otherwise, accept the default, Enable.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.12.1.2

Parameter:	**Mask**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > NSSA Ranges
Default:	None
Options:	Any subnet mask
Function:	Allows you to change the subnet mask you assigned to the range of NSSA type 7 subnets. The network address, together with the mask, specifies the type 7 subnets to be grouped in this NSSA range. Only one type 5 LSA will be generated for all type 7 subnets in this range, rather than one type 5 LSA for each type 7 subnet imported by the NSSA.
Instructions:	Enter the appropriate subnet mask in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.12.1.4

Parameter: Status

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > NSSA Ranges

Default: Advertise

Options: Advertise | Block

Function: Specifies whether the border router advertises type 5 LSAs for an NSSA address range. If you set this parameter to Advertise, the border router advertises a single external route for each NSSA address range. You can hide certain networks from other areas by setting this parameter to Block.

Instructions: Select Block if you want to hide the existence of certain networks from other areas. Otherwise, accept the default, Advertise, to advertise a single route for an NSSA address range.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.12.1.5

Parameter: External Route Tag

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > NSSA Ranges

Default: 0

Options: 0 to 16777215

Function: Specifies the value to insert in the external route tag field of translated type 5 LSAs configured for a type 7 address range.

Instructions: Specify the value (1 through 16777215) that you want to insert in the external route tag field for translated type 5 LSAs. Otherwise, accept the default value, 0.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.12.1.6

OSPF Virtual Interface Parameters

Parameter: Transit Area ID

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces > **Add**

Default: None

Options: Any area ID

Function: Identifies the transit area through which this virtual link is configured.

Instructions: Enter the appropriate area ID in dotted-decimal notation. The transit area must contain the neighboring router identified in the Neighbors Router ID parameter.

MIB Object ID: Not applicable

Parameter: Neighbors Router ID

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces > **Add**

Default: None

Options: Any IP address

Function: Identifies the interface at the other end of this virtual link.

Instructions: Enter the appropriate IP address.

MIB Object ID: Not applicable

Parameter: Enable

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces

Default: Enable

Options: Enable │ Disable

Function: Enables or disables this virtual link. This parameter is useful when you want to temporarily disable a virtual link rather than delete it.

Instructions: Set to Disable to turn off this virtual link. Set to Enable if you previously disabled this virtual link and now want to reenable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.2

Parameter:	**Transit Delay**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces
Default:	1 second
Options:	1 to 360 seconds
Function:	Indicates the estimated number of seconds it takes to transmit a link state update packet over this interface.
Instructions:	Either accept the default value of 1 second or enter a new value from 1 to 360 seconds.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.7.1.6

Parameter:	**Retransmit Interval**
Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces
Default:	5 seconds
Options:	1 to 360 seconds
Function:	Indicates the number of seconds between link-state advertisement retransmissions for adjacencies belonging to this interface. This value is also used when retransmitting database description and link-state request packets. This value should be well over the expected round-trip time. Although the default value is 5, Nortel Networks suggests the following values for this parameter: for broadcast, 10 seconds; for point-to-point, 15 seconds; for NBMA, 15 seconds; for point-to-multipoint, 15 seconds.
Instructions:	Either accept the default value of 5 seconds or set the retransmit interval to some other value from 1 to 360 seconds.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.7.1.7

Parameter: Hello Interval

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces
Default:	15 seconds
Options:	1 to 360 seconds
Function:	Indicates the number of seconds between the Hello packets that the router sends on the interface. Nortel Networks suggests the following values for this parameter: for broadcast, 10 seconds; for point-to-point, 15 seconds; for NBMA, 20 seconds; for point-to-multipoint, 15 seconds.
Instructions:	Either accept the default value of 15 seconds or set the Hello Interval parameter to another value from 1 to 360 seconds. This value must be the same for the virtual neighbor and for all routers attached to the same network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.7.1.8

Parameter: Dead Interval

Path:	Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces
Default:	60 seconds
Options:	1 to 2000 seconds
Function:	Indicates the number of seconds that a router's Hello packets have not been seen before its neighbors declare the router down. This value should be some multiple of the Hello interval. Nortel Networks suggests the following values for this parameter: for broadcast, 40 seconds; for point-to-point, 60 seconds; for NBMA, 80 seconds; for point-to-multipoint, 60 seconds.
Instructions:	Either accept the default value of 60 seconds, or enter another value for this parameter. This value must be the same for all routers attached to the same network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.3.7.1.9

Parameter: **Password**

Path: Configuration Manager > Protocols > IP > OSPF/MOSPF > Virtual Interfaces

Default: None

Options: Any ASCII text string up to eight characters long

Function: Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. This parameter is valid only when the Authentication Type parameter for the area is set to Simplepassword.

Instructions: Enter the appropriate password. All routers in the same area must either have no authentication or have the same password.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.10

RIP Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: Enable

Options: Enable | Disable

Function: Specifies whether the Routing Information Protocol (RIP) is enabled on this interface.

Instructions: Select Enable to enable RIP on this interface. Select Disable to disable RIP on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.2

Parameter: **RIP Supply**

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: Enable

Options: Enable | Disable

Function: Specifies whether the interface transmits periodic RIP updates to neighboring networks.

Instructions: Select Enable to configure the interface to transmit RIP updates. Select Disable to prohibit the interface from transmitting RIP updates.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.5

Parameter: RIP Listen

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: Enable

Options: Enable │ Disable

Function: Specifies whether this interface listens for RIP updates from neighboring networks.

Instructions: Select Enable to configure this interface to listen for RIP updates and, thus, add received routing information to its internal routing table. If you select Enable, a configured policy can still prohibit the interface from updating its internal routing tables. Select Disable to configure the interface to ignore RIP updates from neighboring routers. Thus, the interface does not add received routing information to its internal routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.6

Parameter: Default Route Supply

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: Disable

Options: Enable │ Disable │ Generate

Function: Specifies whether or not the interface advertises a default route in RIP updates sent to neighboring networks. When a router does not know the route of a particular address, it uses the default route as the destination. A configured policy can override this setting. This parameter is independent of the RIP Supply parameter.

Instructions: If you select Enable, RIP advertises the default route if it is present in the routing table -- that is, if you have statically included a default route in the table or if the router has learned the default route (0.0.0.0) dynamically. If you select Generate, RIP advertises a default route whether or not a default route is present in the routing table. (This parameter does not cause RIP to create a routing table entry for a default route; the route will not be visible in the routing table.)

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.7

Parameter: Default Route Listen

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: Disable

Options: Enable │ Disable

Function: Specifies whether or not IP adds default route information to its internal routing table.

Instructions: Select Enable to configure the RIP interface to listen for and potentially add the default route (0.0.0.0) information to its internal routing table. Note that you must also enable RIP Listen on this interface. A configured policy can override this setting. Select Disable to prohibit the RIP interface from adding the default route (0.0.0.0) information to its internal routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.8

Parameter: Poisoned Reverse

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: Poisoned

Options: Poisoned │ Actual │ Split

Function: Specifies how the RIP interface advertises routes it learns from an adjacent network in periodic updates subsequently sent to that network.

Instructions: Select Poisoned to configure this RIP interface to implement poisoned reverse. When poisoned reverse is enabled, the RIP interface advertises routes to the adjacent network from which it learned the routes. In RIP updates, RIP uses a hop count equal to the RIP network diameter plus one, thus declaring the destination unreachable. Poisoned reverse can speed up the convergence of the network routing tables. Select Split to configure this RIP interface to implement split horizon. When split horizon is enabled, the RIP interface omits routes learned from a neighbor in RIP updates subsequently sent to that neighbor. Select Actual to configure this RIP interface to advertise routes with the learned cost. This mode is useful on a frame relay interface that has virtual connections (VCs) to different routers that are part of the same logical IP subnet.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.9

Parameter: Time to Live

Path	Configuration Manager > Protocols > IP > RIP Interfaces
Default:	1
Options:	1 to 255
Function:	Specifies a TTL value to be inserted in the IP header for RIP updates. Certain RIP implementations ignore packets with a TTL value of one hop. Use this parameter to provide interoperability with such implementations.
Instructions:	Setting a TTL of 1 prevents RIP updates from inadvertently exiting the local network. Increasing the TTL introduces the risk of the update exiting the local network and being forwarded around the network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.11

Parameter: Broadcast Timer

Path:	Configuration Manager > Protocols > IP > RIP Interfaces
Default:	30 seconds; 1 hour for dial-optimized routing
Options:	5 seconds to 86,400 seconds (24 hours); 1 hour to 1,209,600 seconds (2 weeks) for dial-optimized routing
Function:	Specifies how frequently RIP generates a full update of the routing table.
Instructions:	Enter a value in 5-second increments.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.12

Parameter: Timeout Timer

Path:	Configuration Manager > Protocols > IP > RIP Interfaces
Default:	90 seconds; 3 hours for dial-optimized routing
Options:	15 seconds to 259,200 seconds (72 hours); 3 hours to 3,628,800 seconds (6 weeks) for dial-optimized routing
Function:	Specifies the time period that RIP will wait for an update for a particular network before declaring it to be unreachable.
Instructions:	Nortel Networks recommends a timeout value of the broadcast time multiplied by 3. Enter a time in 5-second increments.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.13

Parameter: Holddown Timer

Path:	Configuration Manager > Protocols > IP > RIP Interfaces
Default:	90 seconds; 3 hours for dial-optimized routing
Options:	15 seconds to 259,200 seconds (72 hours); 3 hours to 3,628,800 seconds (6 weeks) for dial-optimized routing
Function:	Specifies the time period that unusable routes will be advertised through this interface after the route has become invalid. This parameter affects how long a route remains in the routing table after the route has become unusable. Note that if a route to a destination becomes unusable, the holddown value does not affect the router's ability to learn new routes to the same destination.
Instructions:	Nortel Networks recommends a timeout value of the broadcast time multiplied by 3. Enter a time in 5-second increments.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.14

Parameter: RIP Mode

Path:	Configuration Manager > Protocols > IP > RIP Interfaces
Default:	RIP I
Options:	RIP I | RIP II | RIPII with aggregation
Function:	Specifies which mode of RIP to run.
Instructions:	If you specify RIP I, RIP generates RIP Version 1 packets only. The destination IP address is the directed broadcast address, and the destination MAC address is the broadcast address. Select RIP I if any of the listening devices are RIP Version 1-only devices. If you select RIP II, RIP generates RIP Version 2 updates with the destination MAC address set to the multicast address of 224.0.0.9 specified in the RIP Version 2 RFC. The destination MAC address on Ethernet and FDDI networks will be the corresponding multicast address. On all other media, the destination MAC address will be the broadcast address. RIP does not aggregate subnet information in the updates. If you select RIPII with aggregation, RIP generates Version 2 updates but aggregates subnet information in the manner of RIP Version 1. Nortel Networks recommends using RIP II mode, with or without aggregation, rather than RIP I mode, especially if unnumbered point-to-point links or variable-length subnets are used.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.15

Parameter: Triggered Updates

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: Disable

Options: Enable │ Disable

Function: Configures RIP to generate an update each time it recalculates a route's metric.

Instructions: For compatibility with routers running Version 8.10 or earlier, disable this feature. Implementations of RIP earlier than Version 9.00 do not support triggered updates. If you enable triggered updates, RIP will generate triggered updates with a maximum frequency of one every 5 seconds. The route will include all changes that occurred in the last 5 seconds. This enforced interval prevents RIP from monopolizing CPU resources during periods of instability.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.16

Parameter: Authentication Type

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: None

Options: None │ Simple

Function: Specifies the way RIP handles simple authentication in RIP2 mode.

Instructions: If you are running RIP in RIP2 mode and do not want authentication, set this parameter to None. If you set the parameter to Simple, RIP drops all received Version 1 updates and processes only Version 2 updates with the correct password set.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.17

Parameter: Authentication Password

Path: Configuration Manager > Protocols > IP > RIP Interfaces

Default: None

Options: A valid password string up to 16 characters

Function: Specifies a password.

Instructions: Set the Authentication Type to Simple and enter a password.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.2.2.1.18

Parameter: Initial Stabilization Timer

Path:	Configuration Manager > Protocols > IP > RIP Interfaces
Default:	120 seconds
Options:	0 to 86,400 seconds
Function:	Specifies the interval that RIP uses as its initial stabilization period.
Instructions:	Specify an interval that will allow RIP to learn all routes from its neighbors before sending a full routing update on the interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.19

Parameter: FR/ATM Service Broadcast

Path:	Configuration Manager > Protocols > IP > RIP Interfaces
Default:	Disable
Options:	Enable | Disable
Function:	Specifies whether or not frame relay services are established for all adjacent hosts before sending RIP updates.
Instructions:	Set to Enable if you want frame relay established for all adjacent hosts before sending RIP updates. Otherwise, accept the default, Disable.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.2.2.1.20

Appendix B
Routing Policy Parameters

This appendix contains the Site Manager parameter descriptions for accept and announce policies. You can display the same information using Site Manager online Help.

This appendix contains the following information:

Topic	Page
Accept Policies	
Common Accept Policy Parameters	B-2
RIP Accept Policy Parameters	B-5
OSPF Accept Policy Parameters	B-6
Announce Policies	
Common Announce Policy Parameters	B-7
RIP Announce Policy Parameters	B-17
OSPF Announce Policy Parameters	B-19

Common Accept Policy Parameters

Parameter:	**Enable**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables this policy.
Instructions:	Set to Disable to disable the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.2
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.2

Parameter:	**Name**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies > **Add**
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies > **Add**
Default:	None
Options:	Any alphanumeric character string
Function:	Identifies this accept policy.
Instructions:	Specify a unique name for the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.4
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.4

Parameter: Networks

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies

Default: An empty list

Options: A list of network identifiers

Function: Specifies the networks to which this policy applies. Each entry consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or to a range of networks.

Instructions: Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Use the default empty list to match any route.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.5

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.5

Parameter: Action

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies

Default: Accept

Options: Accept | Ignore

Function: Specifies whether the protocol ignores a route that matches the policy or forwards the route to the routing table manager.

Instructions: Specify Accept to consider the route for insertion in the routing table. To drop the route, specify Ignore.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.6

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.6

Parameter: Route Preference

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies

Default: 1

Options: 1 to 16

Function: Assigns a metric value (the higher the number, the greater the preference) to a route that the protocol forwards to the routing table manager. If confronted with multiple routes to the same destination, the routing table manager may need to use this value to decide which route to insert.

Instructions: Either accept the default value, 1, or enter a new value. Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference and routes for the most specific networks (longest address and mask) should have the highest preference.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.7

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.7

Parameter: Rule Precedence

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies

Default: 0

Options: A metric value

Function: Assigns a metric value to this policy (a policy with a higher value takes precedence over a policy with a lower value).

Instructions: Use this value to specify the order of precedence for policies that match the same route.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.8

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.8

RIP Accept Policy Parameters

Parameter:	**From Gateway**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list.
Instructions:	Use the default empty list to indicate that this policy applies to RIP updates from any router.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.1.1.10

Parameter:	**Received on Interface**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the IP addresses of one or more interfaces on this router. This policy applies to RIP updates received on interfaces that appear on this list.
Instructions:	Use the default empty list to indicate that this policy applies to RIP updates received on any interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.1.1.11

Parameter:	**Apply Subnet Mask**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Accept Policies
Default:	Null
Options:	Null or IP address mask
Function:	Specifies a mask that will override the interface's subnet mask in the presence of networks with variable-length subnet masks.
Instructions:	Supply a mask, set the Action parameter to Accept, and use the default Network parameter (an empty list). If you specify a mask of 0.0.0.0, the router determines which mask to apply. For example, if the network in the update is a subnet of the same network as the receiving interface, the router applies the mask of the receiving interface. If the network in the update is a subnet of a different natural network, the router applies the natural mask of that network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.1.1.12

OSPF Accept Policy Parameters

Parameter:	**Type**
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies
Default:	Any
Options:	Type 1 | Type 2 | Any
Function:	Describes which types of OSPF ASE routes match this policy.
Instructions:	To match either Type 1 or Type 2, use the default, Any.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.3.1.10

Parameter:	**Tag**
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Accept Policies
Default:	An empty list
Options:	A list of tag values
Function:	Specifies OSPF tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain the tag values in this list.
Instructions:	Use the default empty list to indicate that this policy applies to OSPF ASE advertisements with any tag value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.3.1.11

Common Announce Policy Parameters

Parameter:	**Enable**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables this policy.
Instructions:	Set to Disable to disable the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.2
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.2

Parameter: Name

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	None
Options:	Any alphanumeric character string
Function:	Identifies this policy.
Instructions:	Enter a unique name for the policy.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.4
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.4

Parameter: Networks

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of network identifiers
Function:	Specifies which networks will match this policy. Each identifier consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or to a range of networks.
Instructions:	Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Enter an empty list to match any route.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.5
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.5

Parameter: Action

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	Announce
Options:	Announce | Ignore
Function:	Specifies whether or not to advertise a route that matches this policy.
Instructions:	To advertise the route, specify Announce. To drop the route, specify Ignore.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.6
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.6

Parameter: Rule Precedence

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	0
Options:	A metric value
Function:	Specifies a metric value to be used to compare this policy with other policies that a route may match. A policy with a higher metric takes precedence over a policy with a lower metric. In case of a tie, the protocol uses an internal index value assigned to the policy by IP software. (In general, the index value is indicated by the position of the policy in the Site Manager window -- the last policy in the window has the highest index value.)
Instructions:	Use this parameter to assign precedence to policies that match the same route.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.7
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.7

Parameter:	**Route Source**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	Any
Options:	Any | Direct | Static | RIP | OSPF (not valid for OSPF) | EGP | BGP
Function:	Specifies one or more route source identifiers. If you select a route source ID, a route from that source that meets the other criteria of this policy matches the policy.
Instructions:	To specify any source, use the default.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.8
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.8

Parameter:	**Advertise**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of network identifiers
Function:	Specifies network IDs to include in place of the network IDs listed in the route to be advertised.
Instructions:	Specify a non-null value only if the announce Action parameter is Announce. The values you enter in the advertise list determine the action taken. If you supply a list of network IDs, these IDs are advertised instead of the actual IDs in the route. If you use the default (an empty list), the actual IDs are advertised. Note that by default, BGP-4 aggregates subnets into their natural network IDs. If you supply a list that includes the encoding 255.255.255.255/ 255.255.255.255, the actual network IDs are advertised along with the other IDs in the advertise list. This allows advertisement of an aggregate or default along with the actual network. If the actual network is a subnet (and the advertising protocol supports subnet advertisements), the subnet is advertised.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.10
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.10

Parameter: From RIP Gateway

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list, and applies only to RIP-sourced routes when RIP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to RIP updates from any router.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.11
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.11

Parameter: Received on RIP Interface

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the addresses of one or more interfaces on this router. This policy applies to RIP advertisements received on the interfaces in this list, and applies only to RIP-sourced routes when RIP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to RIP updates received on any interface.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.12
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.12

Parameter: From OSPF Router ID

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IDs of one or more OSPF routers. This policy applies to OSPF advertisements authored by a router on this list, and applies only to OSPF-sourced routes when OSPF is included as a route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to OSPF updates from any router.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.13

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.13

Parameter: Received OSPF Type

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: Any

Options: Type 1 │ Type 2 │ External │ Internal │ Any

Function: Specifies which types of OSPF routes match this policy, and applies only to OSPF-sourced routes when OSPF is included as a route source.

Instructions: To match any route type, enter Any. To match any non-ASE route, enter Internal. To match any ASE route, enter External. To match any external type 1 route, enter Type 1. To match any external type 2 route, enter Type 2.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.14

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.14

Parameter: Received OSPF Tag

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: An empty list

Options: A list of tag values

Function: Specifies tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain tag values in this list, and applies only to OSPF-sourced ASE routes when OSPF is included as a route source.

Instructions: Specify one or more tag values. Use the default empty list to indicate that this policy applies to OSPF ASEs with any tag value.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.15

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.15

Parameter: From EGP Peer

Path: Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP address of one or more EGP peers. This policy applies to EGP advertisements authored by a router on this list, and applies only to EGP source routes when EGP is included as a route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to EGP advertisements from any router.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.16

MIB Object ID: OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.16

Parameter: From EGP AS

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of autonomous system numbers
Function:	Specifies one or more autonomous system numbers. This policy applies to EGP advertisements received from EGP peers in an AS on this list and applies only to EGP-sourced routes when EGP is included as a route source.
Instructions:	Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to EGP advertisements from peers in any AS.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.17
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.17

Parameter: Received EGP Gateway

Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the IP address of one or more EGP gateways. This policy applies to EGP advertisements that use a gateway on this list as the next hop, and applies only to EGP-sourced routes when EGP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to EGP advertisements with any gateway address.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.18
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.18

Parameter:	**From BGP Peer**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies the IP address of one or more BGP peers. This policy applies to BGP advertisements authored by a router on this list, and applies only to BGP-sourced routes when BGP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to BGP advertisements from any router.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.19
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.19

Parameter:	**From BGP AS**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of autonomous system numbers
Function:	Specifies one or more autonomous system numbers. This policy applies to BGP advertisements received from BGP peers in an AS on this list, and applies only to BGP-sourced routes when BGP is included as a route source.
Instructions:	Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to BGP advertisements from peers in any AS.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.20
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.20

Parameter:	**Received BGP Next Hop**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies one or more IP addresses. This policy applies to BGP advertisements whose Next Hop attribute matches an IP address on this list and applies only to BGP-sourced routes when BGP is included as a route source.
Instructions:	Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to BGP advertisements with any Next Hop attribute.
MIB Object ID:	RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.21
MIB Object ID:	OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.21

RIP Announce Policy Parameters

Parameter:	**External Route Source**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Default:	Any
Options:	Direct | Static | RIP | OSPF (with Type 2 metric) | EGP | BGP | Any
Function:	Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.
Instructions:	This parameter applies only to OSPF routes that use the new ASE type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.2.1.9

Parameter:	**Outbound Interfaces**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Default:	An empty list
Options:	A list of IP addresses
Function:	Specifies a list of outbound RIP interfaces. If an interface appears in this list, the policy applies to RIP advertisements sent via that interface.
Instructions:	Specify one or more IP addresses. Configure an empty list to indicate that this policy applies to any outbound RIP interface.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.2.1.22

Parameter:	**RIP Metric**
Path:	Configuration Manager > Protocols > IP > Policy Filters > RIP > Announce Policies
Default:	0
Options:	0 or an export metric
Function:	Specifies an optional export RIP metric to use when advertising a route that matches this policy.
Instructions:	Set the Action parameter to Announce. If you use the default, the RIP metric is the routing table metric calculated for RIP plus the interface cost.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.2.1.23

OSPF Announce Policy Parameters

Parameter: OSPF Type

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: 0

Options: Type 1 | Type 2 | 0

Function: Specifies an OSPF ASE metric type to use in advertisements for routes that match this policy.

Instructions: Enter 0 if you want to use the default metric that IP includes in the advertisement, based on the route source. For a BGP, EGP, or RIP route, the default is Type 2. For routes from all other sources, the default is Type 1. Set the Action parameter to Announce.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.22

Parameter: OSPF Tag

Path: Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies

Default: Null

Options: Null or a tag value

Function: Specifies a value for the OSPF external route tag field. If the outgoing route matches this policy, the router places this value in the field.

Instructions: Set the Action parameter to Propagate and set the BGP/OSPF Automatic Tag parameter to Disable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.23

Parameter: BGP/OSPF Automatic Tag

Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	Disable
Options:	Enable | Disable
Function:	Enables BGP/OSPF automatic tag generation.
Instructions:	Select Disable (the default) to use the value you specify with the OSPF Tag parameter. Select Enable to generate a tag according to the criteria in RFC 1403 (or any superseding RFC). This parameter overrides the Tag Generation Method parameter in the Edit OSPF Global Parameters window.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.4.1.24

Parameter: OSPF Metric

Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	0
Options:	0 or an export metric
Function:	Specifies an optional OSPF metric to use when advertising a route that matches this policy.
Instructions:	Set the Action parameter to Announce. If you use the default, the OSPF metric is the routing table metric.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.4.1.25

Parameter: OSPF NSSA Propagate

Path:	Configuration Manager > Protocols > IP > Policy Filters > OSPF > Announce Policies
Default:	Disable
Options:	Enable | Disable
Function:	Specifies whether the propagate bit is set in type 7 LSAs of announced OSPF external routes. If the P-bit is set, the NSSA area border router translates type 7 default route LSAs to type 5 default route LSAs before flooding them to all non-stub areas. If the propagate bit is not set in a particular type 7 LSA, the translator ignores the LSA.
Instructions:	Select Enable to set the propagate bit for all originated type 7 routes that belong to this policy. Otherwise, accept the default, Disable.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.6.4.1.26

Appendix C
Import and Export Route Filters

This appendix contains RIP and OSPF import and export route filter parameters. You can display the same information using Site Manager online Help.

RIP Import Filters

Parameter:	**Import Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default:	0.0.0.0
Options:	Any IP network address
Function:	Identifies, by IP address, the network to which this filter applies. If this field is set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate network address in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.8.1.3

Parameter: **Import Mask**

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: 0.0.0.0

Options: Depends on the address class of the network address

Function: Specifies the range of addresses this filter acts upon.

Instructions: For example, consider Class B network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the subnet ID and the final 8 bits to the host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the network ID portion of the address will be filtered. If you enter the mask 255.255.255.0 for this parameter, the network ID and subnet ID portions of the address will be filtered. If the Import Address field is set to 0.0.0.0, and this parameter is set to 0.0.0.0, then the filter applies to *all* routes. If the Import Address field is set to 0.0.0.0, and this parameter is set to 255.255.255.255, then the filter applies to the *default* route. Enter the mask in dotted-decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.4

Parameter: **RIP Gateway**

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: 0.0.0.0

Options: Any IP address

Function: Identifies, by IP address, the router that is sending the updates. This filter will apply to updates from that router. If this field is set to 0.0.0.0, the filter applies to updates from any router.

Instructions: Enter the appropriate IP address in dotted-decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.7

Parameter:	**Interface**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default:	0.0.0.0
Options:	Any IP address
Function:	Specifies the local IP address of the interface that connects this router to the RIP gateway. This filter will apply only to those updates received on this interface. If set to 0.0.0.0, this filter applies to all interfaces.
Instructions:	Enter the appropriate IP address in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.8.1.8

Parameter:	**Action**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default:	Accept
Options:	Accept | Ignore
Function:	Specifies whether the route is transferred to the routing tables. If this parameter is set to Accept (default), the routing information is sent to the routing tables. If this parameter is set to Ignore, the routing information is dropped.
Instructions:	Either accept the default Accept, or select Ignore.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.8.1.5

Parameter: Preference

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: 1

Options: 1 to 16

Function: Assigns a weighted preference value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, EGP, and RIP.

Instructions: If this hierarchy is acceptable, accept the default value, 1. If you want to grant preference to this RIP-derived route, assign a new preference value from 1 through 16 (the greater the number, the higher the preference). Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference. Routes for the most specific networks (longest address and mask) should have the highest preference. The default preference for static routes is 1, but may be set to any value from 1 through 16. If you want to grant a RIP-derived route preference over a static route, make sure that the preference value you assign to the RIP-derived route is greater than the preference value of the static route that you want it to override.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.6

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters

Default: Enable

Options: Enable | Disable

Function: Enables or disables this import route filter.

Instructions: Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now want to reenable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.2

Parameter:	**Import Apply Mask**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Import Filters
Default:	0.0.0.0
Options:	Any IP address mask
Function:	Specifies a mask that will override the interface's subnet mask in the presence of networks with variable-length subnet masks.
Instructions:	Supply a mask, set the Action parameter to Accept, and use the default Network parameter (an empty list). If you specify a mask of 0.0.0.0, the router determines which mask to apply. For example, if the network in the update is a subnet of the same network as the receiving interface, the router applies the mask of the receiving interface. If the network in the update is a subnet of a different natural network, the router applies the natural mask of that network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.8.1.9

RIP Export Filters

Parameter:	**Export Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default:	0.0.0.0
Options:	Any IP network address
Function:	Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate IP address in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.9.1.3

Parameter: Export Mask

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters

Default: 0.0.0.0

Options: Depends on the address class of the network address

Function: Specifies the range of addresses upon which this filter acts.

Instructions: For example, consider Class B network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the subnet ID, and the final 8 bits to the host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the network ID portion of the address will be filtered. If you enter the mask 255.255.255.0 for this parameter, the network ID and subnet ID portions of the address will be filtered. If you set the Export Address field to 0.0.0.0 and set this parameter to 0.0.0.0, then the filter applies to *all* routes. If you set the Export Address field to 0.0.0.0 and set this parameter to 255.255.255.255, then the filter applies to the *default* route. Enter the appropriate mask in dotted-decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.4

Parameter: From Protocol

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters

Default: Any

Options: Any | RIP | EGP | OSPF | Direct | Static | BGP-3

Function: Identifies the source of the routing information: direct connection, static route, or RIP-, OSPF-, EGP-, or BGP-3-derived route.

Instructions: Select the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.5

Parameter: Interface

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default: 0.0.0.0
Options: Any IP address
Function: Identifies the outgoing IP interface for the RIP update. This filter will apply only to this interface. If set to 0.0.0.0, this filter applies to all interfaces.
Instructions: Enter the appropriate IP address in dotted-decimal notation.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.7

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default: Enable
Options: Enable | Disable
Function: Enables or disables this export route filter.
Instructions: Set to Disable if you want to disable this export route filter. Set to Enable if you previously disabled this export route filter and now want to reenable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.2

Parameter: Action

Path: Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default: Propagate
Options: Propagate | Ignore | Aggregate
Function: Controls the flow of routing information. If you set this parameter to Propagate, this route is advertised. If you set this parameter to Ignore, advertising of this route is suppressed. If you set this parameter to Aggregate, the network is not explicitly advertised. Instead, the default route (0.0.0.0) is advertised.
Instructions: Either accept the default, Propagate, or select Ignore or Aggregate.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.6

Parameter:	**RIP Metric**
Path:	Configuration Manager > Protocols > IP > Route Filters > RIP > Export Filters
Default:	0 (the actual route cost as learned)
Options:	0 to 15
Function:	Assigns a RIP cost to the propagated route. The value 0 causes the actual route cost (as learned) to be used.
Instructions:	Accept the default value, 0, or enter a new value. Do not use a value that exceeds the diameter of the RIP network.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.9.1.8

OSPF Import Filters

Parameter:	**Import Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > OSPF > Import Filters
Default:	None
Options:	Any IP address
Function:	Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate network address in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.10.1.3

Parameter:	**Preference**
Path:	Configuration Manager > Protocols > IP > Route Filters > OSPF > Import Filters
Default:	0
Options:	0 to 16
Function:	Assigns a weighted preference value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, EGP, and RIP.
Instructions:	If this hierarchy is acceptable, accept the default value, 0. If you want to grant preference to this OSPF-derived route, assign a new preference value from 1 through 16 (the greater the number, the higher the preference). Routes for all networks (0.0.0.0/0.0.0.0) should have the lowest preference. Routes for the most specific networks (longest address and mask) should have the highest preference. The default preference for static routes is 0, but it may be set to any value from 0 through 16. If you want to grant an OSPF-derived route preference over a static route, make sure that the preference value you assign to the OSPF-derived route is greater than the preference value of the static route that you want it to override.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.10.1.6

OSPF Export Filters

Parameter:	**Export Address**
Path:	Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters
Default:	0.0.0.0
Options:	Any IP network address
Function:	Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks.
Instructions:	Enter the appropriate IP address in dotted-decimal notation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.11.1.3

Parameter: Export Mask

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: None

Options: Depends on the address class of the network address

Function: Specifies the range of addresses upon which this filter acts.

Instructions: For example, consider Class B network 172.32.0.0. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If you enter 255.255.0.0 for this parameter, only the network ID portion of the address will be filtered. If you enter the mask 255.255.255.0 for this parameter, the network ID and subnet ID portions of the address will be filtered. If you set the Export Address field to 0.0.0.0 and set this parameter to 0.0.0.0, then the filter applies to *all* routes. If you set the Export Address field to 0.0.0.0 and set this parameter to 255.255.255.255, then the filter applies to the *default* route. Enter the appropriate mask in dotted-decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.4

Parameter: Export From Protocol

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters

Default: RIP

Options: Any | RIP | EGP | OSPF | Direct | Static | BGP-3

Function: Identifies the source of the routing information: direct connection, static route, or RIP-, EGP-, OSPF-, or BGP-3-derived route.

Instructions: Select the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.5

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters
Default: Enable
Options: Enable | Disable
Function: Enables or disables this export route filter.
Instructions: Set to Disable if you want to disable this export route filter. Set to Enable if you previously disabled this export route filter and now want to reenable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.2

Parameter: Action

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters
Default: Propagate
Options: Propagate | Ignore
Function: Controls the flow of routing information. If you set this parameter to Propagate, this route is advertised. If you set this parameter to Ignore, advertising of this route is suppressed.
Instructions: Either accept the default, Propagate, or select Ignore.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.6

Parameter: Type

Path: Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters
Default: Type 1
Options: As Is | Type 1 | Type 2
Function: Specifies an OSPF ASE metric type to use in advertisements for routes that match this policy.
Instructions: Select As Is if you want to use the default metric that IP includes in the advertisement, based on the route source. For a BGP, EGP, or RIP route, the default is Type 2. For routes from all other sources, the default is Type 1. Set the Action parameter to Propagate.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.7

Parameter: Tag

Path:	Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters
Default:	0
Options:	0 to 2,147,483,647
Function:	Sets the tag value for the AS external advertisement that is generated for this network. This parameter has meaning only when the Action parameter is set to Propagate.
Instructions:	Enter the appropriate tag.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.11.1.8

Parameter: Auto Tag

Path:	Configuration Manager > Protocols > IP > Route Filters > OSPF > Export Filters
Default:	Disable
Options:	Enable | Disable
Function:	If enabled, the router creates a tag for this route as described in RFC 1364 (superseded by RFC 1403) for BGP/OSPF Interaction.
Instructions:	Set to Enable if you are running BGP-3 as your exterior gateway protocol.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.2.1.11.1.9

Appendix D
Route Weight Worksheet

This appendix provides a worksheet that you can use to calculate route weights for OSPF, RIP, EGP, BGP, and static routes.

1. **Select one route from the following list:**

Direct	0
OSPF internal	0
OSPF external	16 (OSPF import preference)
RIP	16 (RIP import preference)
EGP	16 (EGP import preference)
BGP	16 (BGP import preference)
Static	16 (SR preference)

2. **Multiply the value associated with the route by the following decimal or hexadecimal value:**

   134217728 x _____=_____

   or

   0x8000000 x ____ = _____

3. **Select one route from the following list:**

Direct	0
OSPF intra-area internal	0
OSPF inter-area internal	1
OSPF type 1 external	2
OSPF type 2 external	3
Non-OSPF external (RIP, EGP, BGP)	3
Static	3

4. **Multiply the value associated with the route by the following decimal or hexadecimal value:**

   16777216 x _____ = _____

   or

   0x1000000 x _____ = _____

5. **Select one route from the following list:**

Direct	0
OSPF internal	0
OSPF type 2 external w/ASE metric support enabled	0
EBGP	2
RIP	4
EGP	5
Static	6
OSPF type 2 external w/ASE metric support disabled	7

6. **Multiply the value associated with the route by the following decimal or hexadecimal value:**

   2097152 x _____ = _____

   or

   0x200000 x _____ = _____

7. **Select one route from the following list and calculate the associated value using the formulas supplied:**

Direct	OSPF interface? OSPF interface metric: 0
OSPF imported	OSPF-TOS-0 Metric
BGP-3 imported	Calculate a decimal or hexadecimal value using one of the following formulas: 8192 * (16 - BGP3 Import BGP3 Preference) + (IGP Origin ? 0 : 4096) + (AS Weighted Path Length <= 4095 ? AS Weighted Path Length : 4095) or (0x2000 * (16 - BGP3 Import BGP3 Preference)) + (IGP Origin ? 0 : 0x1000) + (AS Weighted Path Length <= 0x0fff ? AS Weighted Path Length : 0x0fff)
BGP-4 imported	Calculate a decimal or hexadecimal value using one of the following formulas: 8192 * (16 - BGP4 Import BGP4 Preference) + (IGP Origin ? 0 : 4096) + (AS Weighted Path Length <= 4095 ? AS Weighted Path Length : 4095) or 0x2000 * (16 - BGP4 Import BGP4 Preference) + (IGP Origin ? 0 : 0x1000) + (AS Weighted Path Length <= 0x0fff ? AS Weighted Path Length : 0x0fff)

IBGP imported	Calculate a decimal or hexadecimal value using one of the following formulas: 8192 * (16 - BGP4 Import BGP3 Preference) + (Local Pref <= 8191 ? 8191 - Local Pref : 0) or 0x2000 * (16 - BGP4 Import BGP4 Preference) + (Local Pref <= 0x1fff ? 0x1fff - Local Pref : 0)
RIP imported	Metric
EGP imported	Distance
Static	SR cost

8. Add the values you have calculated.

_____

_____

_____

_____

The total is the route weight: _____

Appendix E
IP/OSPF Configuration Example

The IP/OSPF configuration in Figure E-1 shows an AS divided into different types of OSPF areas using several types of OSPF routers and interfaces as well as variable-length subnetting. This appendix describes parameter settings for router 1 through router 6. Routers W through Z and A through D are included in the figure to complete the network map.

In the configuration in Figure E-1:

- The OSPF Area Authentication parameter is set to none for all areas.
- All timer parameters (Hello, dead, and poll interval) are set to their default values.
- No virtual links are configured.
- Route summarization is not used.
- The frame relay cloud is set to the default management type. It is non-fully meshed with all group mode PVCs.
- There are three area border routers (router 2, router 3, and router 4), two internal backbone routers (router 1 and router 5) and one AS boundary router (router 6 has an Ethernet interface using RIP).
- Router 2 is the designated router for the FDDI segment. Router 4 must be the designated router for the frame relay cloud.
- The unnumbered LAN connecting router 6 and router D is included to ensure that every internal node in area 0.0.0.1 is reachable from every other internal node.

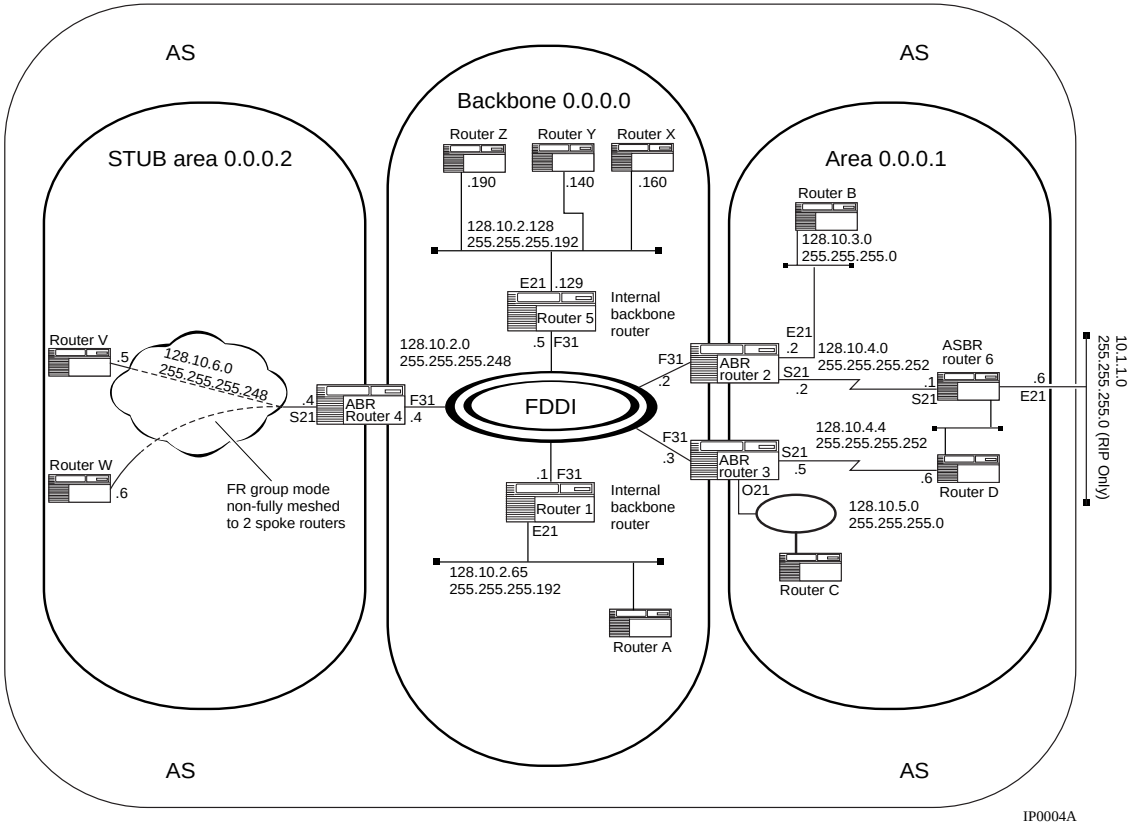

Figure E-1. IP/OSPF Configuration

Tables E-1 through E-6 list nondefault configuration parameters for router 1 through router 6. Parameters that are not shown are set at their default values.

Table E-1. Internal Backbone Router 1

Site Manager Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.1
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.1
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface E21	
IP Configuration/IP Address	128.10.2.65
IP Configuration/Mask	255.255.255.192
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast

Table E-2. Area Border Router 2

Site Manager Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.2
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.2
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
OSPF Interface/Rtr Priority	2 or greater
Interface E21	
IP Configuration/IP Address	128.10.3.2
IP Configuration/Mask	255.255.255.0
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Broadcast
Interface S21	
WAN Protocol	Standard
IP Configuration/IP Address	128.10.4.2
IP Configuration/Mask	255.255.255.252
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Point-to-point

Table E-3. Area Border Router 3

Site Manager Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.3
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.3
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface O21	
IP Configuration/IP Address	128.10.5.3
IP Configuration/Mask	255.255.255.0
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Broadcast
Interface S21	
WAN Protocol	Standard
IP Configuration/IP Address	128.10.4.5
IP Configuration/Mask	255.255.255.252
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Point-to-point

Table E-4. Area Border Router 4

Site Manager Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.4
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.4
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface S21	
WAN Protocol	Frame relay
IP Configuration/IP Address	128.10.6.4
IP Configuration/Mask	255.255.255.248
OSPF Area/Area	0.0.0.2
OSPF Area/AS External	No (area 0.0.0.2 is a stub)
OSPF Interface/Broadcast Type	Point-to-multipoint standard

Note: Router 4 must be the designated router for the frame relay network within the cloud. To ensure this, set the Router Priority parameter on the OSPF frame relay interfaces for router V and router W to 0. The broadcast type should be set to Point-to-multipoint (standard). In addition, router V and router W must have IP adjacent host entries configured for each other.

Table E-5. Internal Backbone Router 5

Site Manager Window/Parameter	Setting
Interface F31	
IP Configuration/IP Address	128.10.2.5
IP Configuration/Mask	255.255.255.248
OSPF Global/Rtr ID	128.10.2.5
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast
Interface E21	
IP Configuration/IP Address	128.10.2.129
IP Configuration/Mask	255.255.255.192
OSPF Area/Area	0.0.0.0
OSPF Interface/Broadcast Type	Broadcast

Table E-6. AS Boundary Router 6

Site Manager Window/Parameter	Setting
Interface S21	
Wan Protocol	Standard
IP Configuration/IP Address	128.10.4.1
IP Configuration/Mask	255.255.255.252
OSPF Global/Rtr ID	128.10.4.1
OSPF Global/ AS Boundary Router	Yes
OSPF Area/Area	0.0.0.1
OSPF Interface/Broadcast Type	Point-to-point
Interface E21	
Add Protocols	RIP
IP Configuration/IP Address	10.1.1.6
IP Configuration/Mask	255.255.255.0

Index

A

accept policies
  OSPF, 6-57
  overview, 1-22
  RIP, 5-24
  setting maximum number, 3-13
acronyms, xix
action command, 6-45
Addr Mask Reply parameter, 3-28, A-11
address
  E.164, 3-33
  IP, 1-2
  IP, for OSPF, 6-4
  MAC, 3-33
  SMDS, 3-38
  WAN, for frame relay network, 3-39
Address Mask parameter, 3-53, A-27
address range, configuring for OSPF NSSA, 6-50
Address Resolution Protocol
  enabling and disabling, 4-2
Address Resolution Protocol (ARP)
  address resolution scheme for, 4-6
  ATMARP, 4-5
  cache timeout feature, 4-11
  customizing global characteristics, 4-3
  datalink encapsulation options for, 4-7
  function of, 1-10
  HP Probe, 4-5
  Inverse ARP, 4-5
  proxy ARP, 4-8
  X.25 DDN and PDN, 4-5
Address Resolution Type parameter, 4-7, A-7
address-resolution command, 4-6
Adjacent Host Address parameter, 3-48, A-30
Adjacent Host Type parameter, 3-48, A-32
Adjacent Host X.121 Address parameter, 3-48, A-32
adjacent hosts, definition, 3-46
aggregate route, definition, 1-8
aging command, 4-11
all-subnet broadcasting, enabling and disabling on IP interface, 3-30
all-subnet-broadcast command, 3-30
all-subnets enabled command, 3-10
announce policies
  OSPF, 6-57
  overview, 1-22
  RIP, 5-24
  setting maximum number, 3-13
Area Address parameter, OSPF, 2-8
area area-id command, 6-40
area command, 6-20
Area ID parameter, OSPF interface, 6-21, A-43
area, OSPF
  description, 1-16
  ID for, 6-20
ARP Forwarding parameter, 4-4, A-20
ARP Server ATM Address Network Prefix parameter, 3-44, A-17
ARP Server ATM Address User Part parameter, 3-44, A-18
ARP. *See* Address Resolution Protocol
arp-buf-limit command, 3-19
arp-mode command, 3-43
arp-server-address command, 3-43
arp-server-reg-interval command, 3-43
AS Boundary Router parameter, 6-8, A-38
ASB parameter, 3-31, A-6

as-boundary-router command, 6-7
as-default-tag command, 6-13
ASE Metric Support parameter, 6-10, A-39
ase-metric-support command, 6-10
ATM (asynchronous transfer mode), IP over, 3-42
ATM ARP Mode parameter, 3-44, A-17
authentication command, 5-7
Authentication Password parameter, 5-7, A-69
Authentication Type parameter
  OSPF area, 6-44, A-51
  RIP interface, 5-7, A-69
authentication-key command, 6-43
authentication-type command, 5-6, 6-43
autonomous system (AS), definition, 1-8

B

backbone, OSPF, definition, 1-16
Backup Enable parameter, 6-6, A-40
Backup Log Mask parameter, 6-17, A-40
backup-log-mask command, 6-16
border router, OSPF, 6-53
bridging, configuring in host-only mode, 3-6
broadcast address
  definition, 3-23
  for subnets, 3-23
Broadcast Address parameter, 3-24, A-5
broadcast command, 3-23
Broadcast Timer parameter, 5-19, A-67
Broadcast Type parameter, 2-8, 3-57, A-33
broadcast-timer command, 5-19

C

cache timeout feature, ARP, 4-11
cache-size command, 3-37, 3-41
circuitless IP interface
  configuring, 2-10
  described, 2-10
  selecting slots for, 2-11
classless interdomain routing (CIDR), 1-8
conventions, text, xviii
cost command, 3-25
Cost for PtP Links parameter, A-53
Cost parameter
  IP interface, 3-25, A-5
  IP static route, 3-53, A-28
customer support, xxii

D

datalink encapsulation options for ARP, 4-7
DDN X.25 address resolution, 4-5
Dead Interval parameter
  OSPF interface, 6-30, A-46
  OSPF virtual interface, 6-56, A-63
dead interval, OSPF, 6-29
dead-interval command, 6-29
Default Route Listen parameter, RIP interface, 5-18, A-66
Default Route Supply parameter, RIP interface, 5-16, A-65
Default TTL parameter, 3-10, A-21
default-listen command, 5-17
default-supply command, 5-16
deleting
  IP from an interface using Site Manager, 3-45
  IP from an interface using the BCC, 3-45
  OSPF from an interface, 6-38
  RARP services, 4-15
  RIP from an interface, 5-23, 6-39
Destination IP Address parameter, 3-53, A-27
dial-optimized routing, 5-18
disabling
  all-subnet broadcasting on IP interface, 3-30
  equal-cost multipath support, 3-15
  global ARP, 4-2
  global IP, 3-3
  ICMP address-mask replies, 3-28
  ICMP redirect messages, 3-29
  IP interface, 3-22
  MTU discovery on an interface, 3-26
  OSPF, 6-3, 6-19
  OSPF area, 6-42

disabling *(continued)*
RARP on an interface, 4-13
RARP on the router, 4-14, 4-15
RIP, 5-3
RIP listening, 5-14
route filter support, 3-14
Router Discovery, 3-55, 3-56, 3-57, 3-58, 3-59
UDP checksum processing, 3-31

E

E.164 address for IP interface, 3-33
ecmp-method command, 3-16
Enable Default Route for Subnets parameter, 3-13, A-24
Enable ISP Mode Support parameter, A-26
Enable parameter
adjacent host, 3-48, A-30
global IP, 3-3, A-19
IP interface, 3-22, A-4
IP static route, A-26
OSPF
area, 6-42, A-51
global, 6-3, A-37
interface, 6-19, A-42
neighbor, 6-38, A-50
NSSA range, A-59
range, A-56
virtual interface, 6-56, A-61
RIP, 5-4, A-64
Router Discovery, 3-55, A-33
Enable Security parameter, A-16
Enable/Disable parameter
RARP global, 4-15, A-35
RARP interface, 4-13, A-35
Enabling
IP interface, 3-22
enabling
all-subnet broadcasting on IP interface, 3-30
all-zero and all-one subnet addresses, 3-10
alternate associated address, 2-13
equal-cost multipath support, 3-15
global ARP, 4-2
global IP, 3-3
ICMP address-mask replies, 3-28
ICMP redirect messages, 3-29
IP interface, 3-22
MTU discovery on an interface, 3-26
OSPF area, 6-42
OSPF boundary function, 6-7
OSPF interface, 6-19
OSPF on the router, 6-3
RARP on the router, 4-14
RIP, 5-3
RIP listening, 5-14
route filter support, 3-14
Router Discovery, 3-55, 3-56, 3-57, 3-58, 3-59
source routing over token ring, 3-34
UDP checksum processing, 3-31
equal-cost multipath, 3-15
Estimated Hosts parameter, 3-12, A-23
Estimated Networks parameter, 3-12, A-23
estimating size of routing table, 3-11
Ethernet Arp Encaps parameter, 4-8, A-12
External Route Tag parameter, 6-52, A-60
external route tag, OSPF, 6-11

F

filters, IP traffic, 1-26
Forward Cache Size parameter, 3-38, 3-42, A-16
forwarding command, 3-4
Forwarding parameter, global IP, 3-5, A-19
forwarding table, maximum size of, 3-40
frame relay network, WAN address for, 3-39
FRM Broadcast parameter, 3-40, A-13
FRM Cast 1 DLCI parameter, 3-40, A-14
FRM Cast 2 DLCI parameter, 3-40, A-14

H

Hello Interval parameter
OSPF interface, 6-28, A-45
OSPF virtual interface, 6-56, A-63
Hello interval, OSPF, 6-27
hello-interval command, 6-28
Hold Down Timer parameter, OSPF, 6-14, A-39
holddown command, global OSPF, 6-14

Holddown Timer parameter, 5-22, A-68
holddown-timer command, 5-21
hops, specifying maximum number with time-to-live value, 3-9
Host Cache parameter, 4-12, A-8
Host Encapsulation parameter, 3-48, A-31
host-only mode
  configuring bridging, 3-6
  configuring global IP for, 3-4
HP Probe, definition, 4-5

I

Import AS Extern parameter, 6-48, A-52
Import Summaries parameter, 6-48, A-53
import-summaries command, 6-47
Initial Stabilization Timer parameter, 5-23, A-70
Interface MTU parameter, 3-37, A-10
Interface Preference parameter, 3-60, A-34
interface, definition, 3-20
Internet Control Message Protocol (ICMP)
  definition, 3-29
  enabling and disabling address mask replies, 3-28
  enabling and disabling redirect messages, 3-29
Internet Network Information Center (NIC), 1-2
Inverse ARP, 4-5
IP
  deleting from an interface using Site Manager, 3-45
  deleting from an interface using the BCC, 3-45
  described, 1-2
  equal-cost multipath, 3-15
  global
    enabling and disabling, 3-3
    forwarding mode, 3-4
  interface
    all-subnet broadcasting on, 3-30
    cost of, 3-25
    E.164 address for, 3-33
    enabling and disabling, 3-22
    MAC address for, 3-33
    MTU discovery on, 3-26
    UDP checksum processing on, 3-31
  policies, maximum number of accept and announce, 3-13
  starting, 2-5
IP address
  acquiring through RARP services, 1-12
  network classes, 1-2
  specifying in dotted-decimal notation, 1-3
IP Address parameter
  IP configuration, 2-5, A-2
  OSPF, 2-8, 2-9
  RARP, 4-14, A-36
  RIP, 2-6
IP OSPF Maximum Path parameter, 3-17, A-25
IP router, internal routing tables, 1-20

L

Lifetime parameter, 3-59, A-34
listen command, 5-14
log-mask command, 6-16

M

MAC Address parameter
  adjacent host, A-31
  IP interface, 3-34, A-9
  RARP, 4-14, A-35
MAC address, for IP interface, 3-33
Mask parameter
  IP interface, 2-11, A-4
  NSSA range, A-58
mask-reply command, 3-28
Maximum Interval parameter, 3-58, A-34
Maximum Policy Rules parameter, 3-14, A-24
metric command, 6-33
Metric Cost parameter, 6-33, A-47
Metric parameter, A-57
Minimum Interval parameter, 3-58, A-33
mode command, 5-10
mtu command, 6-34
MTU Discovery parameter, IP, 3-27, A-6
MTU Mismatch Detect Enable parameter, 6-36

MTU Size parameter, 6-35, A-48
mtu-discovery command, 3-26
mtu-mismatch-detect command, 6-35
multiaccess network, router priority for, 6-23
Multicast Extensions parameter, A-41
Multicast Forwarding parameter, A-48
Multiple Nexthop Calculation Method parameter, A-25

N

Neighbor's IP Address parameter, 6-38, 6-41, A-49
Next Hop Addr parameter, 3-53, A-28
Next Hop Interface Addr parameter, 3-48, A-31
Next Hop Mask parameter, 3-53, A-28
Nonlocal ARP Destination parameter, 4-4, A-21
Nonlocal ARP Source parameter, 4-4, A-20
not-so-stubby areas (NSSA)
  configuring, 6-48
  configuring an address range, 6-50
  definition, 1-17
NSSA Def Route Type parameter, A-55
NSSA Originate Def Route parameter, A-54
NSSA Propagate Def Route parameter, A-54
NSSA Range Mask parameter, A-58
NSSA Range Net parameter, A-58
NSSA Translate parameter, A-53
NSSA. *See* not-so-stubby areas.

O

OSPF
  area ID, 6-20
  backbone, definition, 1-16
  backup soloist, 6-5
  border router, 6-53
  boundary function, 6-7
  configuring cost metrics, 6-31
  database synchronization, 1-13
  dead interval, 6-29
  deleting from an interface, 6-38
  described, 1-13
  enabling and disabling, 6-3, 6-19
  external route tag, 6-11
  Hello interval, 6-27
  IP address for, 6-4
  logging messages, 6-15
  maximum transmission unit size, 6-34
  network type, 6-22
  poll interval, 6-30
  retransmit interval, 6-25
  router priority for multiaccess networks, 6-23
  slot for soloist, 6-5
  specifying a preferred path, 6-31
  starting, 2-8
  summary route, 6-44
  virtual link, 6-54
OSPF areas
  backbone, 1-17
  creating, 6-40
  definition, 1-16
  enabling and disabling, 6-42
  normal (non-stub), 1-17
  not-so-stubby area (NSSA), 1-17, 6-48
  stub, 1-17
OSPF NSSA. *See* not-so-stubby areas.
OSPF route filters, C-11
OSPF Slot parameter, 6-6, A-39
ospf-max-paths command, 3-17

P

Password parameter
  OSPF interface, 6-44, A-47
  OSPF virtual interface, 6-56, A-64
PDN X.25 address resolution, 4-5
Poisoned Reverse parameter, 5-10, A-66
poisoned reverse, RIP updates, 5-9
policies
  definition, 1-22
  OSPF, 6-57
  RIP, 5-24
  setting maximum number, 3-13
policy parameters
  Action (accept), B-3
  Action (announce), B-9
  Advertise (announce), B-11
  Apply Subnet Mask (accept), B-6, C-5

policy parameters *(continued)*
Enable (accept), B-2
Enable (announce), B-7
External Route Source (announce), B-17
From BGP Peer (announce), B-16
From BGP Peer AS (announce), B-16
From EGP Peer (announce), B-14
From Gateway (accept), B-5
From OSPF Router ID (announce), B-13
From RIP Gateway (announce), B-12
Name (accept), B-2
Name (announce), B-8
Networks (accept), B-3
Networks (announce), B-8
OSFP Tag, B-19
OSPF Metric (announce), B-20
OSPF Type (announce), B-19
Outbound Interface (announce), B-18
Precedence (announce), B-9
Received BGP Next Hop (announce), B-17
Received EGP Gateway (announce), B-15
Received on Interface (accept), B-5
Received on RIP Interface (announce), B-12
Received OSPF Tag (announce), B-14
Received OSPF Type (announce), B-13
Route Preference (accept), B-4
Rule Precedence (accept), B-4
Tag (accept), B-7
Type (accept), B-6
Poll Interval parameter, 6-31, A-46
poll interval, OSPF, 6-30
poll-interval command, 6-30
Preference parameter, 3-53, A-29
preference, definition, 1-20
Primary Log Mask parameter, 6-17, A-40
priority command, 6-23
Priority parameter, 6-38, A-50
product support, xxii
Proxy ARP, 4-8
proxy command, 4-9
Proxy parameter, 4-10, A-8
publications
hard copy, xxi
related, xxi

R

Range Mask parameter, 6-46, A-56
Range Net parameter, 6-46, A-55
RARP
customizing, 4-12
defining the mapping table for, 4-13
deleting globally, 4-15
disabling and reenabling on an interface, 4-13
disabling globally, 4-14
overview, 1-12
starting, 2-9
RARP on an interface, 4-13
Redirect parameter, 3-30, A-11
Registration Refresh Interval parameter, 3-44, A-18
Remote Party Sub-Address parameter, 3-48, A-32
retransmission-interval command, 6-26
retransmit interval for OSPF, 6-25
Retransmit Interval parameter
OSPF interface, 6-27, A-45
OSPF virtual interface, 6-56, A-62
Reverse Address Resolution Protocol. *See* RARP
RFC 1583 Compatibility Enable parameter, A-42
RIP
adding to an interface, 2-7
authenticating password on Version 2 update, 5-6
configuring timers, 5-18
default route, 5-16
dial-optimized routing, 5-18
enabling and disabling, 5-3
equal-cost multipath support, 3-15
listening for default route, 5-17
listening for updates, 5-14
poisoned reverse updates, 5-9
sending triggered updates, 5-11
setting diameter, 5-2
split horizon updates, 5-9
stabilization time, 5-22
starting, 2-6
supplying updates, 5-8
time-to-live value for updates, 5-13
update mode, 5-4
RIP Diameter parameter, global IP, 5-2, A-22
RIP Listen parameter, 5-15, A-65

RIP Maximum Equal Cost Paths parameter, 3-17, A-25
RIP Mode parameter, 5-6, A-68
RIP route filters, C-1
RIP Supply parameter, 5-9, A-64
rip-diameter command, 5-2
rip-max-paths command, 3-16
Route Filter Support parameter, 3-15, A-24
route filter support, enabling and disabling, 3-14
route-filters command, 3-14
Router Discovery
  broadcast type for advertisements, 3-56
  definition, 1-20, 3-54
  enabling and disabling, 3-55
  interface preference for, 3-59
  interval between advertisements, 3-57
  lifetime of advertised addresses, 3-58
Router ID parameter, 6-5, A-38
router-id command, 6-4
routing table, estimating size of, 3-11
Rtr Priority parameter, 6-24, A-44

S

size of routing table, estimating, 3-11
Slot Mask parameter, IP interface, 2-11, A-15
slot-mask command, 6-5
SMDS Arp Request Address parameter, 3-39, A-13
SMDS Group Address parameter, IP interface, 3-39, A-12
soloist, OSPF, 6-5
split horizon, RIP updates, 5-9
starting
  IP, 2-5
  IP on circuitless interface, 2-10
  OSPF, 2-8
  RIP, 2-6
state command, 3-55, 3-56, 3-57, 3-58, 3-59
  ARP, 4-2
  global IP, 3-3
  IP interface, 3-22
  OSPF, 6-3
  OSPF area, 6-42
  OSPF interface, 6-19
  RIP, 5-3
static black hole routes
  configuring, A-27, A-28
  definition, 3-53
static routes, definition, 3-49
Status parameter, A-57, A-60
stub areas
  configuring, 6-46
  definition, 1-17
stub command, 6-47
Stub Default Metric parameter, 6-48, A-52
stub-metric command, 6-47
subnet mask
  function, 1-4
  specifying, 1-5
Subnet Mask parameter
  IP configuration, 2-5, A-3
  OSPF, 2-8, 2-9
  RIP, 2-6
subnets
  broadcast address for, 3-23
  definition, 1-4
  enabling all-zero and all-one addresses, 3-10
  unknown, using default route for, 3-12
summary network command, 6-45
summary route, OSPF, 6-44
supernet
  defining black hole for, 3-53
  definition, 1-7
supply command, 5-8
support, Nortel Networks, xxii

T

Tag Generation Method parameter, 6-13, A-41
technical publications, xxi
technical support, xxii
text conventions, xviii
Time to Live parameter, 5-13, A-67
Timeout Timer parameter, 5-21, A-67
Timeout Value parameter, A-41

timeout-timer command, 5-20
timers, configuring for RIP, 5-18
time-to-live command, 3-9, 5-13
timing out entries in the address resolution cache, 4-11
token ring networks, 3-34
TR End Station parameter, 3-36, A-10
TR Endstation ARP Type parameter, 3-36, A-10
traffic filters for IP, 1-26
Transit Delay parameter
  OSPF interface, 6-25, A-44
  OSPF virtual interface, 6-56, A-62
transit delay, estimating for OSPF, 6-24
transit-delay command, 6-24
Transmit Bcast Addr parameter
  IP configuration, 2-5, A-3
  OSPF, 2-8, 2-9
  RIP, 2-6
Triggered Updates parameter, 5-12, A-69
triggered-updates command, 5-11
type command, 6-22
Type parameter, OSPF interface, 6-23, A-43

U

Udp Xsum On parameter, 3-32, A-9
udp-checksum command, 3-31
UnNumbered Assoc Address parameter
  IP configuration, 2-5, A-4
  OSPF, 2-8, 2-9
  RIP, 2-6
Unnumbered Associated Alternate parameter, 2-13, A-17
Unnumbered CCT Name parameter, 3-53, A-29
unnumbered interface, definition, 2-11
update mode, RIP, 5-4

V

version command, 5-5
virtual link, OSPF, 6-54

W

WAN address, configuring for frame relay network, 3-39
weight, definition, 1-21

Z

Zero Subnet Enable parameter, 3-11, A-22